

20 24 30 35 56 62 67



ETI
ELECTRONICS
TODAY INTERNATIONAL

TOMORROW'S
TECHNOLOGY TODAY

6

SUPER
PROJECTS FOR
YOU TO BUILD!

THE ELECTRONIC NOSE

CAN YOUR COMPUTER
TELL THE DIFFERENCE
BETWEEN BEAUJOLAIS
NOUVEAU &
CHATEAU LAFITE?

AN OVERHEATING
ALARM THAT COULD
SAVE YOUR PC

BUILD A HANDY
TRANSISTOR TESTER

ELECTRONIC TOW
GUARD THAT
CAN STOP
CARAVAN THIEVES

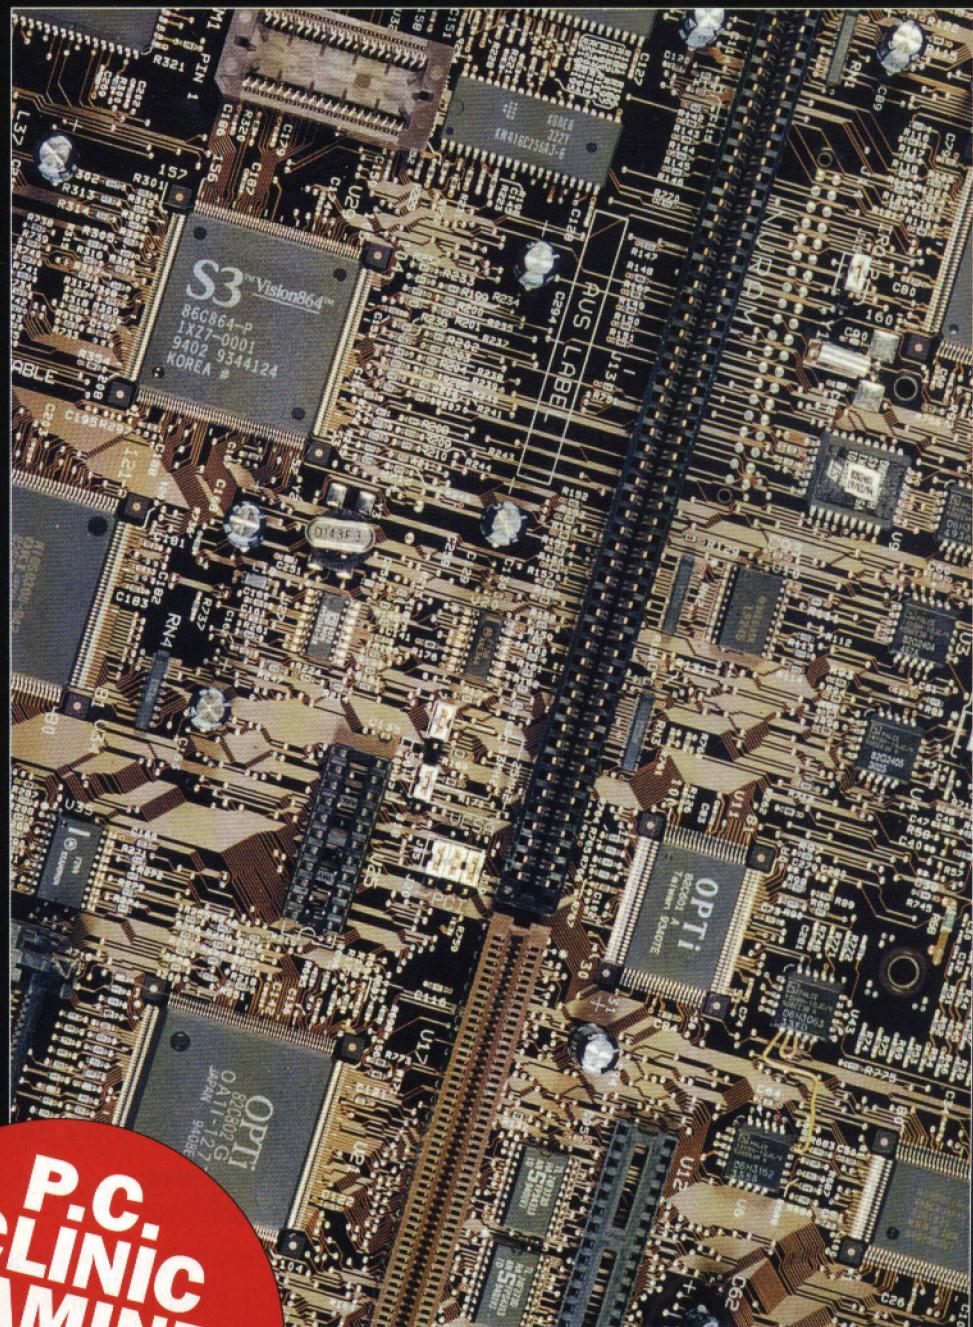
PLUS...
AN INTRODUCTION TO MIDI

COMPUTER CONTROLLED
STEPPER MOTORS

BUILDING A FLUXGATE
MAGNETOMETER

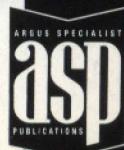
AND...

A LOOK AT ETI'S
NEW
TRANSPUTER
PROJECT

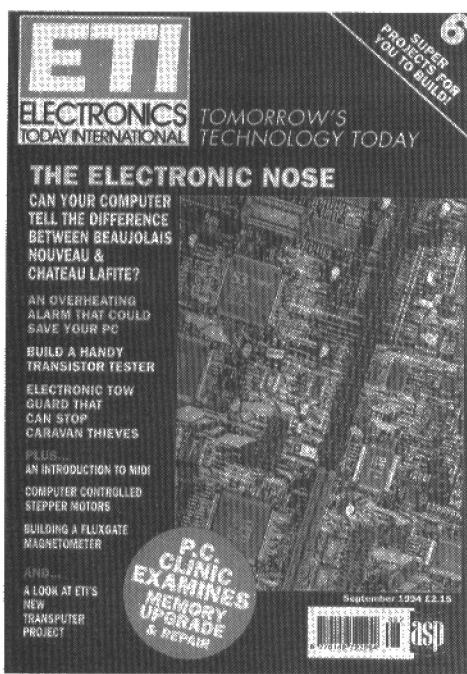


P.C.
CLINIC
EXAMINES
MEMORY
UPGRADE
& REPAIR

September 1994 £2.15



Contents



Volume 23 No.8

Features & Projects

The Electronic Nose 12

A look at advances in bioelectronics, the fusion of electronics and biochemistry, which have already produced a range of sophisticated sensors that can rival the human nose, but in the future could bring us computers so small that they are made from a single molecule.

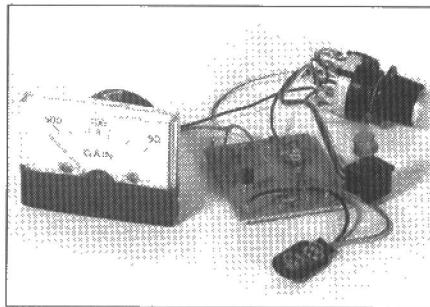
PC Overheating Alarm 20

If your PC overheats because the fan is not working properly or the environment prevents proper cooling, then you could be in serious trouble. To prevent this happening, Robert Penfold has developed a temperature alarm that will plug into your PC and warn you if it is starting to overheat.

The ETI Transputer 24

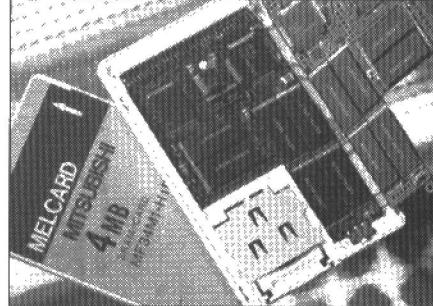
The start of what is probably the world's first practical project to build a single board computer system based around the famous Inmos Transputer chip. It can be used in a wide variety of applications from a graphics engine, to artificial intelligence, or could even be used to build a parallel super computer.

Transistor Tester 30



This simple constant current transistor tester, designed by Robert Penfold, will enable you to easily check whether those surplus transistors in your component box actually work.

PC Clinic 35



Part 4 of the series which shows readers how to repair, maintain, upgrade and build circuits for personal computers. In this issue we look at the main system memory, how it is used and organised and how it can be upgraded.

SUBSCRIPTION & BACK ISSUES HOTLINE
0737 768611

Subscribe & Save

Phone the hotline and take advantage of our special offer detailed on page 68

Tow Bar Alarm 44

Thieves will think twice about trying to steal your caravan if it is fitted with Terry Balbirnie's alarm system

A Touch of Magnetism 48

In Part 2 of this two part article, Keith Garwell looks at the design of sensitive instrumentation which can detect minute changes in the Earth's magnetic field.

An Introduction to MIDI 56

In Part 4 of this series Robert Penfold continues his introduction to the mysteries of MIDI - the now universally accepted standard for communications between electronic musical instruments.

Microprocessor Motor Control 62

Jim Spence concludes his project to use the ETI Experimenter's Computer to accurately control a number of stepper motors.

Regulars

News and event diary 6

Practically Speaking 67

The new series in which Terry Balbirnie divulges some practical hints and tips for the electronics enthusiast.

PCB foils 69

Open Forum 74

Front cover photograph courtesy of IBM, shows the company's new OEM Systemboard Series all-in-one board for 486 CPU up to DX4/100 Local Bus Video (64-bit) and Local Bus IDE, PCI and VESA.

GUIDED MISSILE CONTROL CABLE

4,200 mtrs

only £14.99

UNBELIEVABLE CABLE BARGAIN



SUITABLE FOR UP TO THE MINUTE ELECTRONICS ENTHUSIASTS
this top quality, government standard exceptionally thin 4 core cable.

This cable is *almost impossible to detect with the naked eye* when used for domestic applications.

A mega 4,200 meters on each reel, very thin four core cable, originally manufactured for the SWINGFIRE and other missile systems. 4 insulated wires 0.12mm dia of copper plated high tensile steel overall dia approximately 0.5mm. Each core is insulated with Lewmex T Enamel and overall the four cores are lapped with count cotton and wax impregnated. Supplied on a plastic reel 6.25" x 5.5" and weighs only 2.75kg.

Almost too many suggested uses, security wiring, communications, remote control, binding, garden, farm, invisible security wiring, plant support, bird scaring, wire rope for models and school mechanics, site layouts, fabric and plastic reinforcement, alarm wiring for whole walls, and of course for fishing with its substantial 28lb breaking strain!!!

NEWS FLASH
NEWS FLASH
NEWS FLASH

MIDLANDS STORE
NOW OPEN IN
WOLVERHAMPTON

NEWS FLASH
NEWS FLASH
NEWS FLASH

BULL ELECTRICAL
T/AS "WALTONS"
55a Worcester Street
Wolverhampton
Tel: 0902 22039

Monday - Saturday
9.00am - 5.30pm

(SORRY NO MAIL ORDER SERVICE AVAILABLE FROM
THIS OUTLET)

BULL ELECTRICAL

250 PORTLAND ROAD HOVE
SUSSEX BN3 5QT

MAIL ORDER TERMS: CREDIT CARD, CASH, PO OR CHEQUE
WITH ORDER PLUS £3.00 P&P (UK only) plus VAT.

Please allow 10-14 days for delivery

TELEPHONE ORDERS WELCOME

TEL: 0273 203500

FAX: 0273 323077

ELECTRONICS TODAY INTERNATIONAL



MODELS

Wired remotes for boats and planes



COMMUNICATIONS

Field phones/Buzzers and telephones



GARDEN

Bird scaring, trip wires and plant support



FARM

Fence blinding, Crop support, alarms



ALARMS

Connect up systems and conceal under wallpaper, carpet, etc



SCHOOL

Physics and Science Lab Art Department



REINFORCEMENT

For fabric, plastic and window displays



FISHING

Illuminate floats with multi colour LED's



WIRING

Give Models a realistic look with working mini bulbs. Cars with flashing lights, aeroplanes, dolls houses with room lights, etc.

PRIORITY ORDER FORM

PLEASE RUSH ME AN ORDER OF MISSLE CABLE BARGAIN BY RETURN POST:

NAME:

ADDRESS:

POSTCODE.....

QTY: REEL(S) @ £14.99 PER REEL: TOTAL: £.....

PLEASE SEND ME YOUR FREE
64 PAGE BUMPER CATALOGUE

ADD P&P (per order) £3.00:

I AM AN EXISTING CUSTOMER

PLUS V.A.T. @ 17.5% :

TOTAL :

PAYMENT METHOD: CHEQUE/CREDIT CARD/P/O

CREDIT CARD NO:

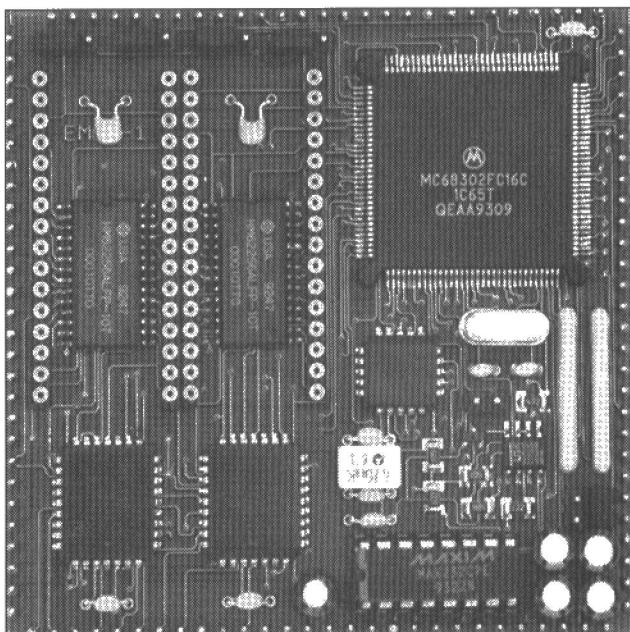
EXPIRY DATE:

68000 Embedded Controller

Developed for OEMs who need the 16-bit power of the 68000 processor in a small module, the EM68 features the highly integrated MC68302 together with 1MB of EPROM, 1MB of 5V FLASH EPROM and 64KB of static RAM in a module measuring just 3in square.

The EM68 has 3 full high speed serial ports operating in UART, HDLC/SDLC, BISYNC or DDCMP modes, with on board RS232 on channel 1, DMA, Interrupt controller, 28 parallel I/O lines, 2 16-bit Timers with compare and capture, Watchdog Timer, DRAM refresh, and Low power (standby) modes.

Designed as a high level component for OEMs to mount directly on their PCB and priced accordingly, the EM68 can cost as little as £95 each in quantity. Editors, Assembler and C compilers are also readily available. For further details contact the manufacturers, Diss, Norfolk based Devantech Ltd, on 0379 644285.

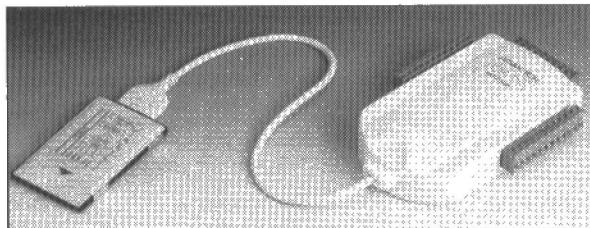


PCMCIA Data Acquisition Pod

Southampton based Integrated Measurement Systems have just launched a PCMCIA data acquisition card, the DACpad. This has been designed to PCMCIA Type II standard and is ideal for laptops, notebooks and host computers with PCMCIA slots, thus allowing such systems to be used for data acquisition in spite of the PC-Bus expansion slot limitations which have in the past prevented them from being used in such applications.

The DACpad provides 8 channels of differential analogue input measurement, with 12-bit resolution at sampling rates up to 30K samples/second. The DACpad includes a signal wiring and conditioning pod, with removable connectors, enabling transducer signals to be wired to the DACpad very simply and easily. On board signal conditioning and amplification allows a wide variety of sensors to be wired directly to the DACpad, including thermocouple CJC. Four separate digital inputs and four digital outputs are also provided for digital input sensing as well as output on/off control.

For further details contact IMS on 0703 771143.

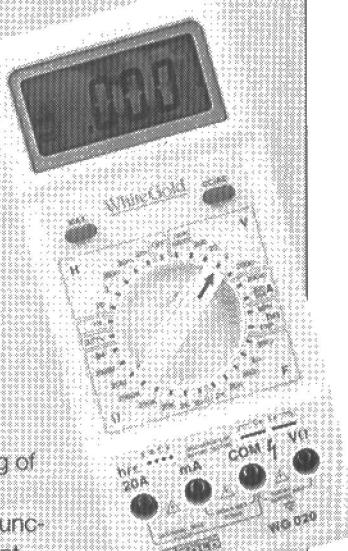


Low Cost Digital Multimeter

New from Maplin Electronics is the WG020 Digital Multimeter. This is a quality, versatile multimeter that provides a wide range of useful features in one, easy to use package.

This highly accurate multimeter features a comprehensive range of functions for the serious hobbyist and professional alike, at an affordable price. The frequency function is autoranging and a 'MAX' hold facility allows the highest reading of a varying measurement to be stored. In addition to the basic functions, a 'duty cycle' measurement facility is provided, which is a useful aid for analysing pulsed signals in digital circuits.

Readings on all ranges are sampled at a rate of 2.5 times per second. Supplied with a pair of test leads, battery and full instruction manual. The WG020 costs £54.95 and for further details contact Maplin Electronics on 0702 554161.

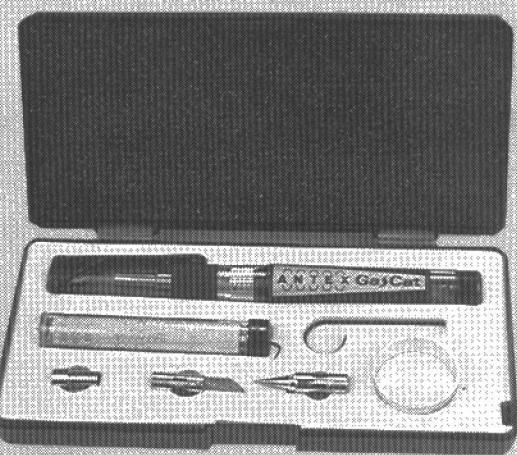


Soldering Irons go Portable

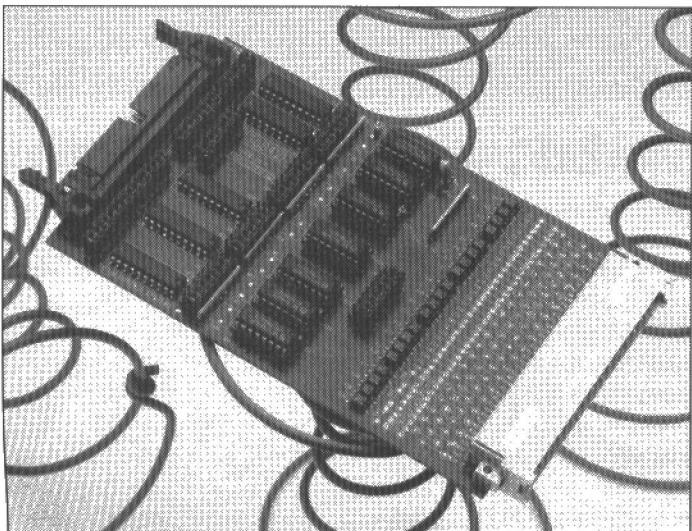
Anyone who is faced with the likelihood of having to do some soldering without a handy mains socket into which to plug a soldering iron will find this new portable gas soldering iron kit from Maplin Electronics an essential item in their toolbox.

This is a compact, handy, pocket sized soldering iron kit which works on butane gas (standard lighter fuel). The highly cost effective kit contains an iron with a 1mm tip, a tube of 1mm 60/40 resin cored solder, a cleaning pad, a hot flame tip, a hot knife tip and a spring clip that fits into two holes as an iron rest when the unit is being used.

The whole kit is supplied in a plastic case with full instructions, spare tips are available separately. This gas soldering iron kit costs £32.95 and for further details contact Maplin Electronics on 0702 554161.



High Density Signal Conditioning Board



A 40 channel opto isolated digital input board has been introduced by Arcom Control Systems of Cambridge. This board provides designers with a single high density module capable of being connected to STE-, VME- and PC/PCAT bus based control and monitoring systems. Designated the SCB42, this single Eurocard sized module simplifies system integration and enhancing operational reliability.

The 40 channels are arranged as five independent ports of eight. Input voltages can range from 0V up to 12V or 24V, with a further user option to fit resistors for inputs up to 48V. Inputs can be of either polarity thanks to the use of opto-isolator devices with back to back LEDs. PCB tracking and component layout have been carefully designed to ensure that a full 500V isolation is maintained across the entire module. Digital debounce circuitry is provided on each channel, with the length of delay determined by the value of a single oscillator capacitor.

For further details contact Arcom on 0223 411200.

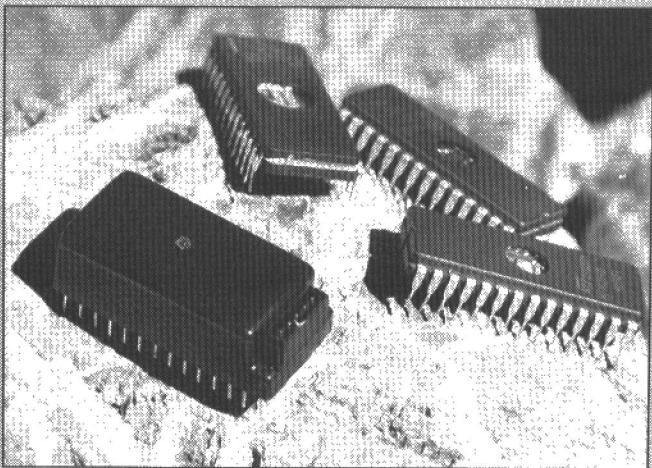
Revolutionary EPROM Emulator

The MicroROM is a revolutionary new type of EPROM emulator developed by London based Squarewave Electronics. The emulator is a module just 11mm high which plugs directly into the EPROM socket and is compatible with standard EPROMs. There is no need to remove, erase and reprogram an EPROM each time data is changed.

To use the emulator, simply plug the module into the EPROM socket, then send the data directly from any computer with a Centronics printer port via a special cable. This cable plugs into a special connector on the emulator module and one can download the entire contents of a 27256 in about 1 second.

Once programming is complete and the connector removed, the module can remain in place like a normal EPROM due to its non volatile storage capability. This makes the MicroROM a lot easier to use and a lot more versatile than conventional emulators.

For further details contact Squarewave on 081 880 9889.



Tomorrow's Technology Update

In the July issue of ETI we carried a feature on smartcards and forecast that they were something we would all be using within two or three years. Judging by the announcements of new developments and new applications which have appeared in the last couple of months, the smartcard revolution might be upon us sooner than we thought.

As an update on the use of smartcards for public transport ticketing systems, the London Transport world leading scheme based in Harrow is forging ahead and has proved very popular with passengers of all ages. At the time of writing, the number of cards issued has just topped the 13,000 mark. The success of the London Transport smartcard project has prompted a large number of other transport operators around the world to start their own schemes for using contactless smartcards.

In May, Manchester started a trial on 120 buses in Bolton and have so far issued about 3,300 cards to children and elderly passengers. Hong Kong are expected to announce a contract shortly for a major system on their Mass Transit Railway and bus services. Meanwhile Japanese railways are starting trials in several stations in Tokyo (several groups of Japanese engineers and developers have come to London over the last few months to see how the London Transport smartcard system is working). Other smartcard ticketing schemes are now being planned in Melbourne, Paris, Oslo, and Helsinki.

In other areas smartcard technology is also advancing rapidly. One of the world's leading RISC processor manufacturers, the UK based company ARM (jointly owned by Acorn Computers, semiconductor manufacturer VLSI and US computer giant Apple) has announced that it is teaming up with a group of European companies under the European Community EITC initiative to develop smart card security systems incorporating analysis techniques such as voice recognition.

Called CASCADE, these new smartcards will be designed around ARM's 32-bit RISC processor technology to produce a card that has over a hundred times the processing power of any existing smartcard. This processing power will be used in conjunction with ARM's high level programming language support and its ability to handle 32-bit data to handle prob-

High Power DSP Prototyping System

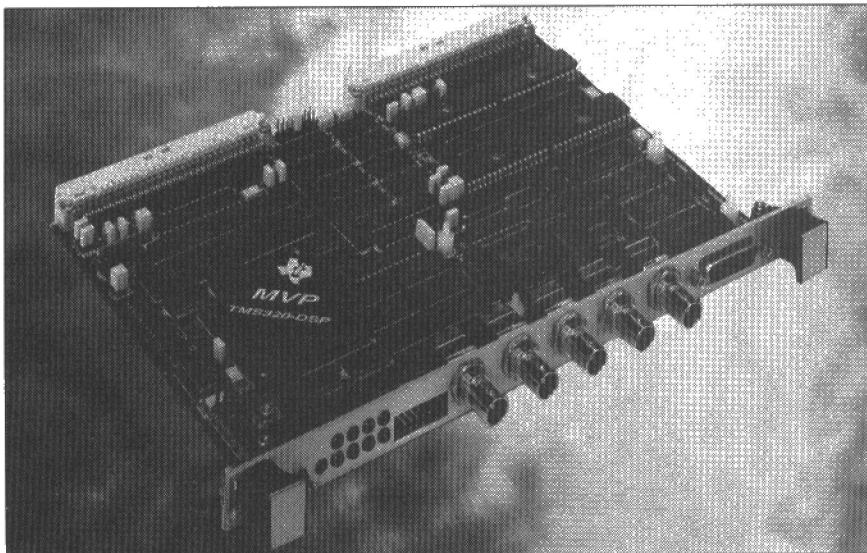
The world's first development and prototyping system for Texas Instruments new high power MVP, the TMS320C80 digital signal processor, has been launched by Loughborough based Loughborough Sound Images (this company is in fact the world's largest supplier of board level products based on DSPs). It's a product which should allow users to rapidly develop and prototype video, image processing and graphics applications.

The development environment, known as EVM, includes general purpose configurable hardware, a range of interface modules for targeted applications and software drivers for the hardware interfaces. These various components will enable engineers to rapidly prototype MVP based designs and cut development costs and time to market.

It is expected that the EVM will be used by designers of MVP based products for video conferencing, document image processing, graphics acceleration and virtual reality systems. The development system can be tailored to each application by connecting interface modules onto the EVM board.

The EVM board includes an MVP processor, integrated 4MB video frame buffer, integrated audio interface and extensive DRAM and synchronous SRAM. Two expansion interfaces, the video interface module and applications interface module allow the addition of application specific plug in modules for hardware flexibility. Later this year, LSI will have video conferencing, virtual reality, and graphics/X-Terminal emulation modules available.

For further details contact LSI on 0509 231843.



lems involving complex algorithms such as voice recognition. This will allow smartcards to reach previously unattainable levels of security and protection against fraud.

In yet another application, Schlumberger has won contracts to develop a smartcard system for the French Health Service. Called CPS the cards will be issued in 1996 to over 1 million healthcare professionals, doctors, nurses, dentists, etc., and is intended to be part of a move aimed at substantially improving operational efficiency.

CPS will provide users with a way of exploiting the power of personal and portable computer systems with a highly secure way of verifying the user's identity, to allow the authorisation of computer based documents and transactions or access to protected data - locally or over telephone/computer networks. This smartcard system should thus help to overcome most of the objections to the use of computerised patient records, since it will effectively deny access to anyone who is not authorised.

There is also a follow-up to the feature on Global Positioning Systems in the August issue of ETI. Bournemouth based Marilake Instruments, a world leading specialist in the design of passenger information systems, has in conjunction with the computing and electronics departments at Bournemouth University, developed an in-flight passenger information system which incorporates a GPS.

Called the Flight Master System it is designed to use the new generation of aircraft seat back video monitors to display a range of maps, flight data and other information for passengers, automatically suspending operation during in flight video, but superimposing essential safety graphics. The built in GPS will enable passengers to see exactly where the plane is at any time during the flight and associated maps will make it easy to identify geographical features, mountains, towns, rivers, etc., which can be seen from the windows.

Contacts:

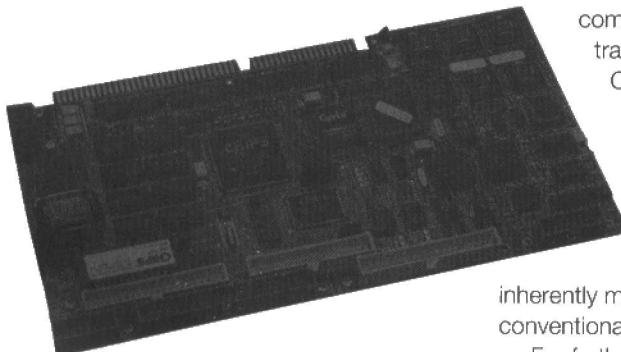
London Transport: 071 918 4123

ARM : 0223 400400

Schlumberger of France: (33) 1 47 46 70 20

Marilake: 0202 570055

Bus Based PC Processor Board



Arcom Control Systems of Cambridge has released a bus-based PC-AT compatible processor board optimised for data acquisition and data concentration applications. Designed around the powerful 50 or 25 MHz 486SLC CPU, with up to 4MB of dynamic RAM, 1.5MB of battery backed static RAM, and 3MB of FLASH EPROM, the board's six serial communications gateways provide an ideal foundation for a wide variety of communications and networking oriented embedded PC systems.

The 486SLC is constructed on a compact PC AT bus module. This approach enables an extremely high performance 486 class industrial PC to be constructed on a passive backplane, providing an inherently modular system with far greater configuration and upgrade flexibility than conventional motherboards - and is ideal for OEM applications.

For further details contact Arcom on 0223 411200.

Event diary

2 August

Talk on aerials, Sudbury and District Radio Amateurs. Tel: 0787 313212

7 August

Radio Society of Great Britain, Annual spares sale Woburn Abbey, Woburn. Tel: 0525 290666

21 August

Southend and District Radio Society Radio and Computer Rally, Rocheway Centre, Rochford, Essex.

Tel: 0702 353676

3 September

Wight Wireless Rally, National Wireless Museum, Arreton Manor, near Newport, Isle of Wight.

Tel: 0983 567665

4 September

Applied Optics and Optoelectronics Conference, Institute of Physics, York. Tel: 071 235 6111

5 September

9th Electromagnetic Compatibility Conference, Institute of Electrical Engineers, Armitage Centre, Manchester. Tel: 071 240 1871

13-15 September

Virtual Reality User Show 1994, London Novotel. Tel: 081 995 3632

20-25 September

Live 94, The Consumer Electronics Show, Earls Court, London. Tel: 0891 500 103

13 Nov

Midland Amateur Radio Society rally at Stockland Green Leisure Centre, Slade Road,

Erdington, Birmingham. Doors open at 10AM, admission £1, For further details ring 021 422 9787 or 021 443 1189 (evenings only).

If you are organising an event which you would like to have included in this section please send full details to: ETI, Argus House, Boundary Way, Hemel Hempstead, Herts. HP2 7ST. Clearly marking your envelope Event Diary.

Bioelectronics

From Biosensors *To* Molecular Computers

Bioelectronics is the name given to the rapidly developing techniques which seek to combine advanced biochemistry and semiconductors, to create a range of highly sophisticated devices that will undoubtedly form the basis of the next great leap forward in technology.

For almost forty years, electronics engineers have been reducing the size of circuits. Valves gave way to transistors, transistors to ICs, and simple gate logic ICs to complete systems on a single chip. Today we can use processor chips with over 3 million transistors and memory chips capable of storing 4MB are commonplace.

The relentless pressure is still there to pack more circuitry on a single chip and to make it operate at higher speeds. To do this, engineers have had to devise techniques for building smaller and smaller transistors, connectors and other components on a silicon chip. A process which is getting increasingly difficult - indeed, many experts consider that practical limits will have been reached within the next decade.

The problem is that conventional electronics are fabricated in what is essentially a 'top-down' process. In other words a large scale design is miniaturised to such a degree that it can be put onto a single silicon chip. The limitations arise from the precision with which the components can be etched onto the chip and already conventional photoetching techniques are coming up against limits imposed by the wavelength of light. Limitations also arise as a result of microscopic flaws in the semiconductor material increasing chip failure rates.

The fact that limitations would eventually be reached was recognised a long time ago and has led some researchers to look for alternative ways of creating ultra-small electronic circuitry. In particular, they have been examining a number of 'bottom-up' approaches which seek to build a circuit molecule by molecule, rather than etching ever finer circuits from a slice of semiconductor.

Some researchers are seeking to build devices from conductors and semiconductors that rely on techniques such as

quantum effects, by using equipment like the tunnelling electron microscope which allows very small structures to be built by moving individual atoms. Other researchers, influenced by the fine detailed structures found in biology, structures such as the DNA molecule, are seeking to build a new generation of devices that are based primarily upon organic molecules rather than traditional metal/semiconductor structures. These are the so called bioelectronic, or molecular electronic devices.

Molecular Electronics

The ultimate goal of researchers into molecular electronic devices, or MED as they are often referred to, is the creation of independent functional units which consist of single molecules. This would mean that a lump of such molecular electronic 'matter', the size of a sugar cube, would contain tens of millions of independently functional circuits that in theory could be as complex as a computer. Circuits which can exist on their own and which may or may not interact with other circuits.

The fact that true MEDs are independent, molecular sized functional units is important, since it differentiates them from compounds which display electrical activity as a result of large numbers of molecules acting in bulk rather than individually. A typical example of an organic molecule which shows bulk electrical properties is the liquid crystal used in LCD displays. At the moment, the majority of the molecular electronic developments with commercial applications are for compounds which show bulk electronic properties.

Perhaps the first person to seriously think about the development of molecular sized machines was the physicist Richard Feynman in about 1959. He realised that molecular machines could be constructed that would store, manipulate and transfer information and he also realised that what made them so

special was not just their microscopic size but the fact that they would be self assembling, self organising and self replicating. In other words, they would share many of the attributes of living organisms.

Feynman's ideas remained little more than theoretical until the mid 1980's, when major advances in genetics and biochemistry, plus the computer simulation of chemical processes and molecular structures, made it possible for researchers to start practically investigating the idea. Pioneer researchers such as Ari Aviram at IBM, a scientist who had been toying with the ideas for such molecular devices since the early 70s, and Robert Birge at Syracuse University. Today there are people working on molecular electronics all over the world and in the UK there are teams at several universities, including Cambridge, Birmingham, UMIST and Warwick.

These researchers have already done a lot of work on developing molecules which could be used as the building blocks for more sophisticated devices. We have thus seen the development by a number of teams of molecular wires, molecular switches, molecular light sources and sensors, plus a very wide variety of molecular sensors. As any electronics engineer will realise, molecular switches are a key

component of any future MED, since they will enable the construction of logic gates and thus of molecules, capable of computational functions. The molecular wires will be needed to connect the switches into functional circuits, using self assembly techniques that have been developed using methods first devised by genetic engineers.

A diverse range of molecular wires and switches have been demonstrated, not all of which rely upon switching a current flow on and off. Some use photons of a specific 'colour' light to cause the molecule to switch from one form to another, a sort of molecular optical computing device. Others use wavelike disturbances of molecular energy, called solitons, to transmit information. The molecules themselves range from pure organic compounds to ones that incorporate metal atoms.

However, one thing is already obvious, we should not expect MEDs to work or behave like anything else we have developed, certainly not like conventional electronic devices.

Molecular Electronic Assemblies

Because of their minute size, it will be impossible to assemble MEDs with conventional techniques, one will certainly not be able to solder wires to them! This means that special techniques have to be developed to both interface such devices to conventional electronic circuitry and to actually build the devices in the first place.

The way in which they are interfaced will depend upon the kind of switching technique which is employed. Where the switching involves an electrical current, then the interfacing technique will involve chemically attaching molecular wires to special field effect transistors which are part of a fairly conven-

tional silicon IC chip. The molecular wires will then be attached to the actual molecular device. This type of interface is already being used, in a very simple manner, in biosensors, where the different types of organic sensor molecules are attached to a semiconductor sensor array with its associated signal amplifiers, interface and reading logic.

One great advantage of optical MED switches is that they will not need any direct connection. Switches such as those being developed by Robert Birge at Syracuse are based upon bacterial rhodopsin and a pulse of laser light at a specific wavelength causes it to switch from one form to another, light of another wavelength will cause it to switch back. The state of switches such as this can be read using another laser beam. Since this

type of switch is both very fast and very small it could well find future applications in MED memory systems.

MEDs are of course three dimensional structures, unlike conventional ICs which are two dimensional. They will, however, in most cases be built upon some form of substrate, with the molecule being built up using self assembling properties borrowed from biological structures such as viruses. Indeed, some researchers are even proposing the use of genetically engineered cells to create MEDs, just as infected cells turn out

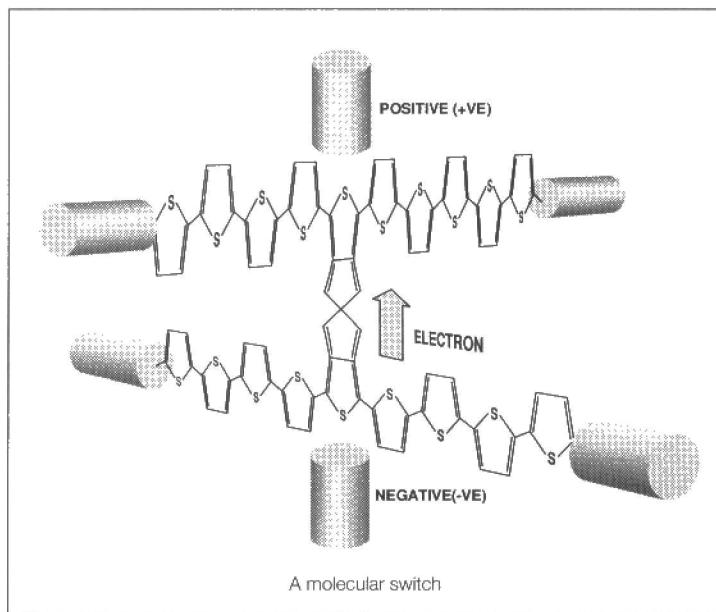
viruses, although it should be noted that such proposals are pure speculation and current knowledge is insufficient to be able to do it.

One technique for organising functional molecules into a practically useful structure is to trap the molecules in a thin film. The technology for developing such films is already well known, for example bilayer lipid membranes are just two molecules thick and Langmuir-Blodgett films are just one molecule thick. Both types of film have been successfully employed to organise bioelectronic molecules in an orderly structure. One example of this technique is a biosensor developed by the Ajinomoto Company of Japan.

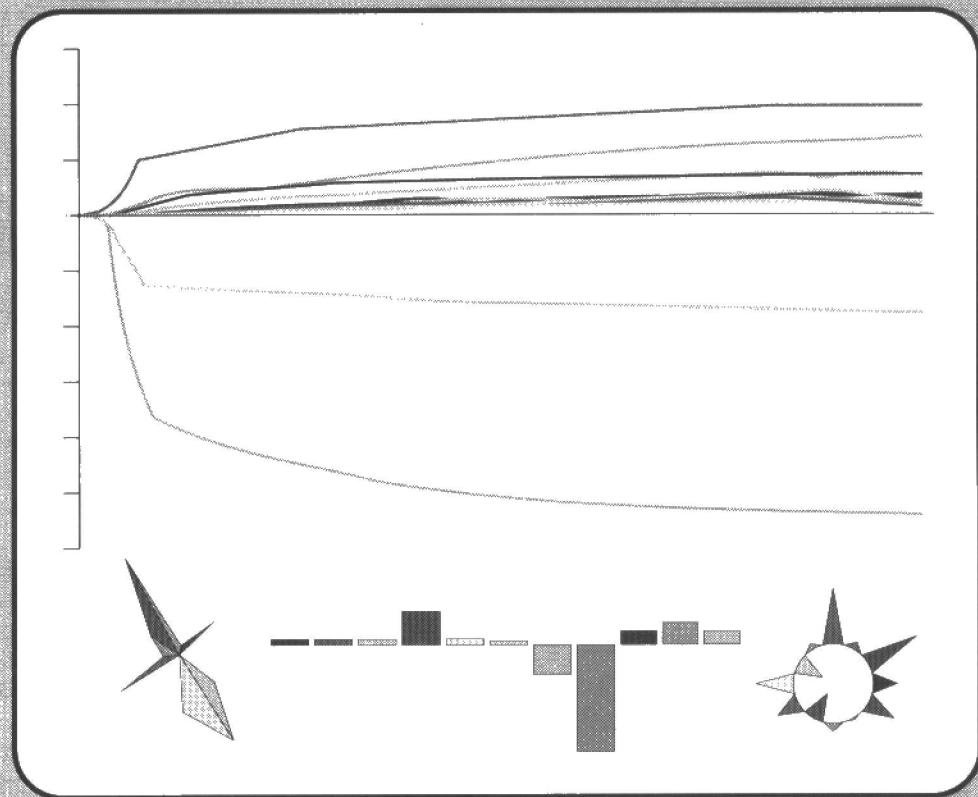
This sensor uses bacteriorhodopsin molecules trapped in a lipid bilayer and the sensor emits electrical signals when struck by light of a specific wavelength. In conjunction with semiconductors, this type of optical biosensor could eventually lead to the creation of highly advanced imaging systems.

Such molecular scale films are being employed in many other similar applications and offer enormous possibilities for future development since films can be stacked up layer on layer to create complex systems. Indeed, stacked molecular scale films offer the best immediate approach for practical construction of practical commercial MEDs.

Slightly thicker, but still extremely thin, films are also being used to take advantage of the bulk electronic properties of organic molecules. A large group of organic molecules are capable of conducting electricity to various extents, some even exhibiting semiconducting properties. These are the type of molecules which have been proposed for use as the wires connecting MEDs. Thin films of such molecules are increasingly



The Electronic Nose



Electronic nose screen display - champagne or white wine?

being used in light weight miniature devices such as batteries, as a replacement for metal conductors. Thus, the newly developed lithium/polymer batteries offer some of the best power to weight ratios thanks to the use of conductive polymer films replacing heavy and bulky metal electrodes.

Indeed, organic polymer films are an area of molecular electronics where a lot of exciting developments are being made, developments which are being turned into practical commercial products right now. One example is the development by scientists at the Cambridge University of an electroluminescent polymer which offers higher efficiencies and brightness than conventional light emitting technology.

The Cambridge team used polyphenylenevinylene, or PPV, as the light emitting polymer. Very thin films of highly purified PPV will emit large amounts of light - indeed, it is about 10% efficient compared to only 1% for an LED. In its natural form PPV emits a green light, but the researchers have managed to produce derivatives which output light in a wide range of colours including both red and a very good blue. The high efficiency of such films means that they produce light that is ten times as bright as that generated by a CRT tube. Early applications will thus be in areas such as emergency lighting and signs. The polymer also offers high reliability and stability with a 100,000 hour life expected.

Biosensors

Another area where organic polymers are already being used commercially with considerable success is in the production of so called biosensors. These are electronic devices which are

capable of detecting very small quantities of a wide range of different chemicals, devices which look set to revolutionise medical and environmental diagnostics. Sensors which, as the title of this article suggests, could provide the computer with a sense of taste and smell that might soon rival and surpass our own.

Already an enormous range of different biosensors have been produced both commercially and experimentally, but all share the same common structure. They use one or more sensor proteins, usually an enzyme, to generate a signal and then a conventional transducer to convert that signal into an electronic signal that can be measured by conventional circuitry and instrumentation.

Thus we can build a biosensor that will measure the amount of glucose sugar in a liquid by using the enzyme glucose oxidase. This was one of the first biosensors to be invented and, in fact, dates back to work done in the 1950s by Leyland C Clark Jr at Cincinnati Children's Hospital. Clark's sensor measured the amount of dissolved oxygen in a solution using a standard platinum electrode and reference electrode, with a membrane permeable to gases. The voltage bias of the platinum electrode was set so that the rate of current flow through the circuit depended on the rate at which oxygen diffused through the membrane, which in turn was proportional to the oxygen concentration in the solution.

In the early 1960s, Clark extended his oxygen sensor to the measurement of glucose concentrations by coating the oxygen sensor with a layer of gel containing glucose oxidase and then, over that, a semi-permeable dialysis membrane that allowed

We greatly underrate our sense of taste and smell. It is truly amazing how an individual, such as a wine expert with a trained nose and palate, can perform what amounts to a very sophisticated chemical analysis on an unknown glass of wine and then say with often unerring accuracy what the wine is, where it came from, what year it was made and even what field the grapes grew in.

To try and produce an electronic 'nose' which can compete with even the limited capability of the average human might seem rather optimistic, but this is exactly what the Hertfordshire based company Neotronics Technology have succeeded in doing. Neotronics became involved in the design of the electronic nose to characterise beer in 1990. It was then part of a DTI funded programme with Bass Brewers and Warwick University.

They have continued and further developed on this initial work and have produced a sensor system which is designed to 'sniff' out the chemical composition of foods, drink, perfumes, household products and, dare it be said, even wine. A sensor system which should help businesses to produce consistently high quality products, to detect counterfeit products and to monitor waste products for potentially dangerous pollutants.

The major advantage of an electronic system over a human nose is that its output is consistent, it can be used 24 hours a day, seven days per week and it will produce numerical output which can be processed by a computer to produce accurate comparative data. There is also the fact that few people will want to spend their working lives sniffing and identifying potentially objectionable smells.

The human nose works by trapping smells in the nasal cavity that lies above the roof of the mouth. At the top of the nasal cavity within the olfactory bulbs there are about 10,000 special non-specific chemical sensors which

perform the task of identifying certain odours. The output of these sensors goes directly to the brain where they are compared with previous odour memories and a conclusion drawn as to the nature and origin of the current odour.

The Neotronics electronic nose aims to do exactly the same as the human nose, but with fewer sensors and a simulated brain. In fact, the current version of this system uses just twelve sensors, but within specific areas, for example wine, it can successfully replicate the function of the human nose.

The twelve sensors are mounted as four sets of three on the inside of a square metal framework which is attached to a cap, allowing a seal to be made between the vapour to be tested and the outside environment. The sensors are connected to control electronics and a computer. The samples to be tested are contained in a glass beaker to which a magnetic stirrer or heater can be added if necessary.

The sensors used in the electronic nose are a special design which has been patented by Neotronics. They are known as conducting polymer sensors and are sensitive to a wide range of different substances. Two different types of polymer are used, polypyrrole and polyaniline. Variations between sensor materials are due to the use of different counterions and solvents which modify the polymer chain.

The conductive polymers are electrochemically grown at a defined potential across an electrode gap. The current design uses two gold electrodes, with the polymer grown in a defined window area, the rest of the sensor being suitably protected. A silicon substrate is now being used for greater reproducibility. A range of sensor materials are currently available and new materials are being developed.

These conductive polymer sensors are very sensitive to certain vapours and odours but, like the sensors in the

glucose molecules to diffuse into the sensor but prevented the enzyme from diffusing out. The enzyme bakes down glucose molecules into gluconic acid and hydrogen peroxide, with the latter creating a signal at the oxygen electrode. The signal at the oxygen electrode is thus proportional to the concentration of glucose molecules in the solution.

The next great step forward in the production of biosensors was based upon the work of James Angell of Stanford and Kensall Wise of Michigan Universities, and involved the marriage of semiconductors and biosensors. They built multiple miniature electrodes on a silicon chip to make electrochemical

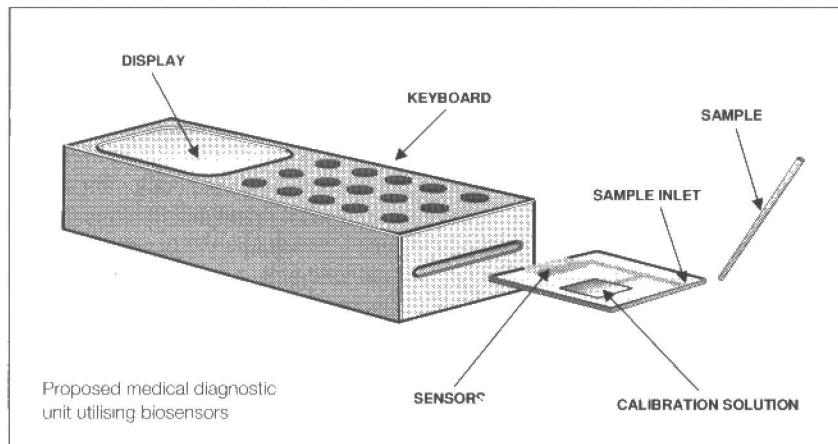
measurements of neural tissue. This technique was extended by Jiri Janata of the University of Utah, who coated the gate of a field effect transistor with the antibody concanavalin A, thus creating the first CHEMFET detector.

This early work was followed by a whole range of different transducers including thermistors, piezoelectric and optical

devices. It also led to considerable work on miniaturisation of sensors and sensor arrays, work that was largely fuelled and funded by the military need for efficient chemical and biological warfare sensors that could give early warning and analysis of an attack, using nerve gas or biological toxins.

The miniaturisation and fusion of biosensors with semiconductors means that whereas Clark's glucose sensor was about a centimetre in diameter and needed a box of electronics, it is now possible to produce a sensor array on a single chip of silicon, with sensors measuring just a few hundredths of a millimetre and with all the necessary transducer and amplifier electronics

mounted on the same chip. To produce such sensors the chemical reagents and membranes are laid down in precise patterns on the semiconductor substrate using an ink jet type printer and the accuracy is such that arrays can be produced that can measure a number of different chemical compounds at the same time.

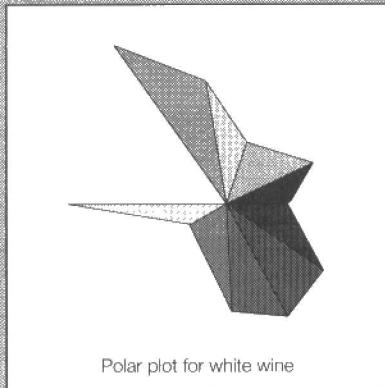


human nose, are not selective, so cannot be used in isolation. In operation, the resistance of the sensor varies in accordance with the concentration of certain molecules within the vapour. And, like the human nose it has an initial sharp response followed by a slower long term effect, which can be seen in the diagrams accompanying this piece.

This change in sensor resistance can be attributed to several different mechanisms. Examples are direct absorption of the vapour into the polymer, modification of the polymer chain by the vapour due to the polarity of the vapour, or the filling of 'holes' within the material by the gas or vapour. There is still a great deal of academic argument as to the exact reason why polymer sensors work, but they do work and that is enough to allow production of a commercial device like the Neotronics nose.

Because conductive polymer sensors are very sensitive to low levels of certain gases and vapours, but not selective, they need to be used in conjunction with sophisticated analysis techniques. In other words, they need to be attached to a computer, via special purpose interface circuitry, on which is running some form of analysis software.

The simplest form of analysis technique employed by Neotronics is to have the computer output the data in graphical form as a scaled polar plot with the maximum vector length defined. This visual picture of a particular 'smell' is easy for the user to



Polar plot for white wine

The development of biosensors has now reached a stage where they can be produced cheaply enough so that a medical worker, for example, can throw away a sensor after each patient measurement. This opens up the possibility of doctors, or even the patients themselves, making instant measurements of the body's physiology. Indeed, such a sensor placed within a patient's vein could be used to continuously monitor blood chemistry in the same way that other equipment can measure heart rate, brain function, etc.

The real power of biosensor technology is that it can provide an instant analysis without any need for a laboratory or skilled technicians. Thus a diabetic can use a biosensor chip to accurately determine when an insulin injection is necessary, as with such a device it becomes much easier to take frequent measurements.

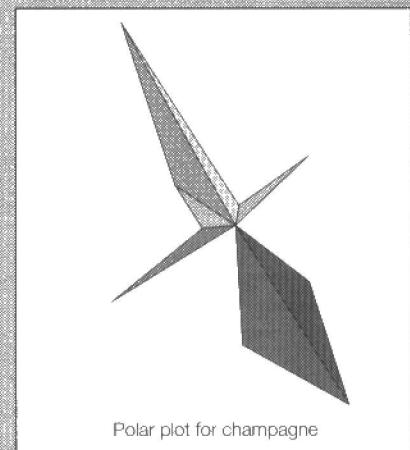
A Japanese company has already developed an implantable artificial pancreas which contains a biosensor for measuring blood glucose levels. A microprocessor interprets the measurements and a small pump automatically gives the patient accurate amounts of insulin that are adjusted to match his sugar metabolism on a minute by minute basis.

Ease of use in the field is an important factor in favour of biosensors in environmental testing, as it can give the user an instant measurement of pollution levels without the time consuming necessity of collecting and analysing samples in the lab. In many cases such measurements can only be made in

interpret, particularly when colours are used. On these polar plots, such as the examples accompanying this piece, each vector represents the output from one sensor, so with twelve sensors there are twelve vectors. Therefore as the relative response from each sensor varies by the addition of a gas or vapour, the overall shape of the polar plot also varies.

This means that each specific substance has its own unique plot shape, in other words a unique signature and a skilled user can recognise familiar substances from their plot. From this was developed the difference ring that allows two different samples to be compared. One would probably be a known reference sample and the other the sample under test. Here, the circle represents no deviation between the two signatures. For each sensor, a vector is taken from the radius, if it goes inward then it is a negative change and if it goes outward a positive change. The radius is scaled so that absolute deviation can be shown.

Once again a skilled user can understand the graphical output and detect any serious deviation from the expected norm. But having to have a skilled user is a major drawback to the usability of such systems. Unfortunately, we cannot simply use the computer to perform this task by comparing patterns with libraries of previously tested samples. This is because



Polar plot for champagne

the field, measurements such as biological oxygen demand which determines the health of rivers and lakes. Speed is of the essence in environmental testing, the quicker pollution can be detected the quicker it can be stopped, the cheaper it is to clean up and the less damage it does.

The same goes for the food industry, where biosensors can be used to keep a constant check on product quality by testing at frequent intervals for certain key chemicals thus reducing the likelihood of having to discard large faulty batches. Already, biosensors are in wide use in the brewing industry to ensure consistency and uniformity of product.

In the future, as can be seen from the box entitled 'Electronic Nose', biosensors will become ever more sophisticated. Soon there will hardly be a chemical substance which it will not be possible to build a biosensor to detect. Furthermore, the detection process will become increasingly sophisticated as a result of the merging of biosensor arrays and sophisticated computer techniques, such as the neural network which has been added to Neotronic's Nose to enable it to learn to detect the subtle differences between a range of different substances. In a couple of years they could well have a system that could, for example, identify an unknown wine with the same unerring accuracy now exhibited by a handful of Masters of Wine.

Biocomputers

Biosensors are the main commercial driving force in bioelec-

every sample of the same substance will vary slightly, the result, amongst other things, of random molecular movement, of temperature and humidity changes.

The solution is a simulated brain, that can learn about the vapours and gases that the electronic nose attached to it encounters. A simulated brain that, like the real thing, can recognise a pattern despite slight variations between the sample and the reference pattern. This is in fact none other than our old friend the neural network.

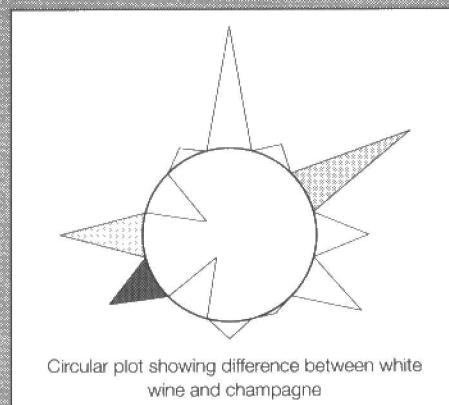
A neural network, either implemented in hardware or software, is essentially a pattern recogniser. It is so good at being a pattern recogniser that it will still correctly recognise a pattern when parts of it are missing or distorted. It also has the added advantage that a neural network is also an efficient learning system. These factors make a neural network the ideal analysis technique for something like the electronic nose with its non-selective sensors.

With a neural network attached to the sensor array electronics we have a system which can be trained to recognise specific substances. The user no longer needs to know what the input signature for those substances look like and neither do variations in individual sensors matter (the signature on one 'nose' system for a specific substance will not look like the signature for that same substance on another

tronics today, but a great many researchers in both commercial and university labs all around the world are working on a much longer range project, the molecular computer. As was mentioned at the beginning of this piece, building very small electronic devices can only be achieved using a bottom up approach to building them. Furthermore, it is perfectly reasonable to conjecture that if it is possible to build a molecular switch and a molecular wire, then in time we should be able to put together molecular circuits of greater and greater complexity.

The ultimate goal for the creation of increasingly complex molecular circuits is of course the construction of a molecular information processing system, a molecular computer. Increasingly, scientists are coming to the conclusion that it will be possible to build such devices, the question being when rather than if. It could be a very long time before molecular computers become a possibility, probably fifty years or more. But on the other hand the rate of technological development has increased enormously over the last few decades and looks set to continue increasing. In this case we could see the molecular computer as a reality in the early decades of the next century.

This brings us to the question of why anyone should want to build computers that are so small that millions of them would fit into a teaspoon. One thing is certain - they will not be used to make small PCs, neither will the actual molecular computers be



Circular plot showing difference between white wine and champagne

system). Neither will slight variations in input signature matter, as once it has been trained it will repeatedly recognise that substance whenever it is encountered.

Of course the neural net does not actually know anything about the substances under test, it will not know anything, for example, about champagne and white wine. But as it is trained by being repeatedly exposed to reference samples, it will build up pattern weightings that will depend on how often they are activated whilst being trained. Once trained, it will recognise that substance when it is encountered, even if it has been watered down, or disguised (such as an attempt to disguise the presence of drugs or explosives). It will even distinguish between the real thing, such as a perfume and a very carefully crafted fake.

This combination of the electronic nose and the neural network has been made possible by a joint development programme between Neotronics and another British high technology company, Neural Technologies. The result of this collaboration is a twelve sensor electronic nose with a neural controller chip. This looks as if it will find a considerable number of uses within the food and drinks industry and, with the development of portable units, could be an added weapon in the battle against terrorists and drug smugglers.

In the not too distant future, a couple of years maybe, we are likely to see much more sophisticated systems with more sensors and capable of recognising a far greater number of different substances. Then it will certainly be practical to train such a system to identify a glass of wine, and tell us the vineyard and vintage and yet another human skill succumbs to technology!

Contacts:

Neotronics - 0279 870182

Neural Technologies - 0730 260256

programmable in high level languages. The molecular computer will probably be an extremely simple computational device with the program hard coded into its structure, its strengths will lie in its enormous parallelism and the way that individual units connect together.

In fact, it is unlikely that molecular computers will be used in any way that even closely resembles the way that computers are used today. Instead, it is likely that they will form part of the structure of things. Scientists have hypothesised that they could be used to make intelligent materials, things that could change shape, colour, surface properties and even repair themselves on the basis of data and energy input. Applications for such intelligent material would be legion.

Molecular computers would also find wide ranging applications in robotics and medical prosthesis. They could form the basis of synthetic muscles. They could form a direct interface between the human brain and machines and allow the construction of systems that would allow the blind to see and the deaf to hear as well as any normal person. And, perhaps they just might form the basis on which we can construct machines that can think for themselves.

This may all sound like science fiction and in many ways it is, but it is science fiction which stands a good chance of becoming science fact within the lifetime of many of us, and perhaps a lot sooner.

PC TEMPERATURE ALARM

If your PC overheats it could damage expensive components but this project from Robert Penfold will warn you when you are getting warm.

Over the years the complexity of PCs has increased substantially. Modern PCs have more complex processors, more memory and more sophisticated graphics cards than the original PCs.

Paradoxically, the amount of power consumed by the average PC has actually reduced slightly.

Modern technology has given us computers that have most of the circuitry on just a handful of chips and reduced power consumption is a by product of this LSI and VLSI Approach

Some PCs still have power supplies that are capable of supplying around 200W, but apparently few PCs actually consume anything approaching this figure.

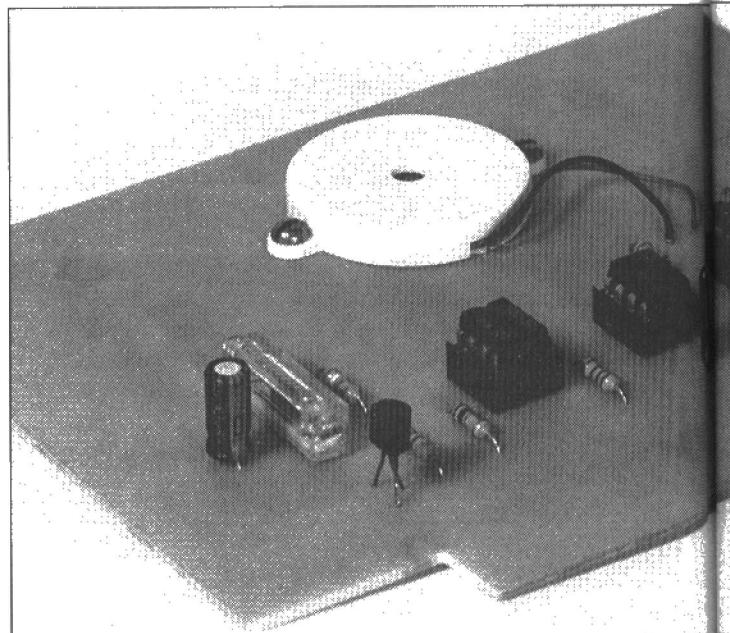
On the other hand, apart from some portable and small desktop computers which use the latest micro-power components, most PCs still consume significant amounts of power and generate a certain amount of heat.

The temperature inside the average PC starts to rise well above the ambient temperature soon after switch-on. Some of the larger integrated circuits run quite hot and if the temperature inside the PC rises too far, there is a risk that these devices will not be able to lose heat at a high enough rate. This could in turn lead to them failing before too long.

Various means of combating overheating are available and these range from simple temperature alarms to up-market devices such as temperature activated fans to keep microprocessors cool.

The simple device featured here is a temperature alarm which activates an audio 'beeper' if the temperature inside the PC exceeds a preset threshold temperature. This temperature is user adjustable and can be anything from 0 to 100 degrees Centigrade.

The unit is in the form of a small PC expansion card which



is simply plugged into any available slot in the host PC. The unit is powered from the PC and its current consumption is only about 12mA

How It Works

Figure 1 shows the block diagram for the PC temperature alarm. The basis of the unit is a semiconductor temperature sensor. All semiconductors are affected by temperature changes and it is normally a potential cause of problems. A semiconductor temperature sensor puts this phenomenon to

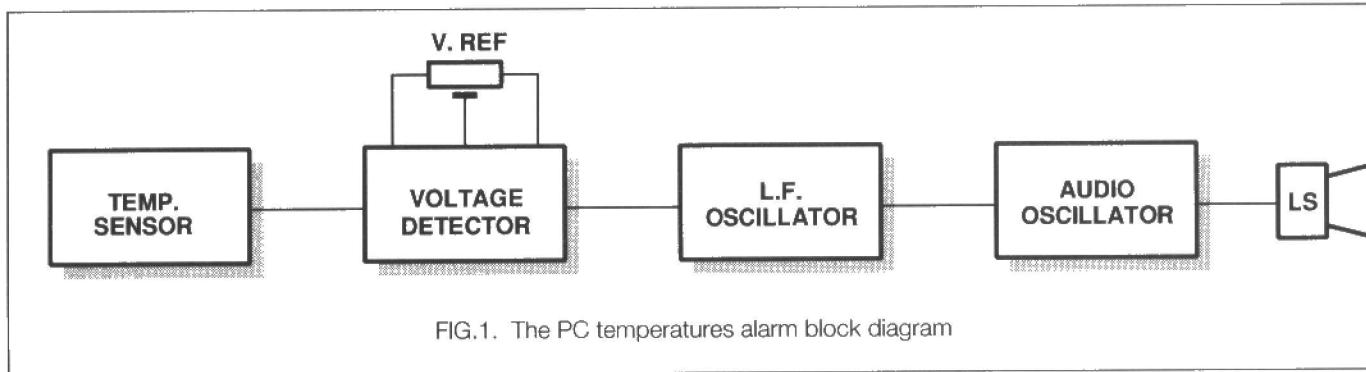
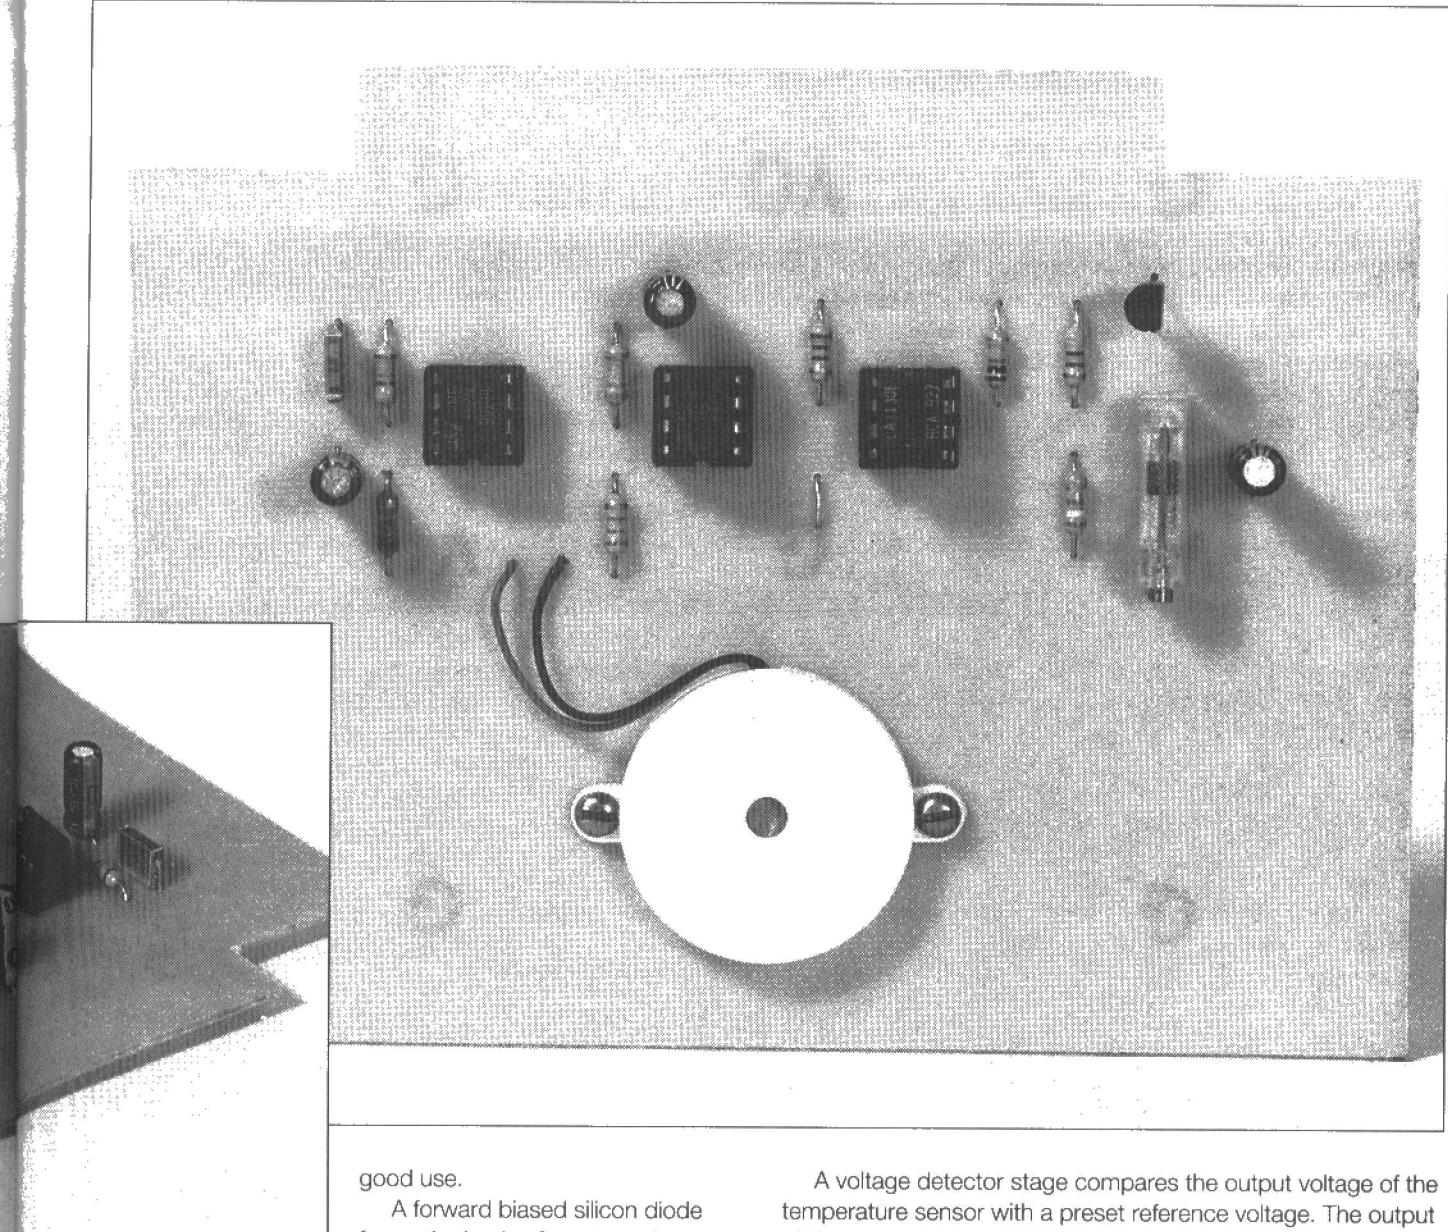


FIG.1. The PC temperatures alarm block diagram



good use.

A forward biased silicon diode forms the basis of most semiconductor temperature sensors. There

is a potential of about 0.65V developed across a forward biased silicon diode, but the exact voltage is dependent to some extent on the bias current used and the temperature of the diode. If the diode is biased with a constant current, the output voltage is only dependent on the diode's temperature.

Even quite large changes in temperature do not provide major changes in the output voltage. The change is typically about two or three millivolts per degree Centigrade, with positive changes in temperature giving a reduction in the output voltage.

The linearity is quite good however, which makes semiconductor temperature sensors a good choice for temperature measurement, as well as simply detecting a certain temperature. In order to use a diode sensor in most practical applications, it is necessary to use a certain amount of amplification, plus some level shifting or other signal conditioning.

The sensor used in this project provides a substantial amount of on-chip signal conditioning, including amplification, level shifting and a phase inversion. As a result of this it, provides an output potential that is equal to 10mV per degree Centigrade. It therefore provides an output voltage of 0 to 1V over its 0 to 100 degree operating range.

A voltage detector stage compares the output voltage of the temperature sensor with a preset reference voltage. The output of the comparator goes high if the output potential from the sensor exceeds the reference voltage. When this happens, the voltage comparator switches on a low frequency oscillator, which in turn pulses an audio oscillator on and off. This generates a simple 'beep - beep' alarm sound, which is adequate for an application of this type. The reference voltage determines the temperature at which the alarm is activated.

The required reference voltage (in millivolts) is equal to the required threshold temperature multiplied by ten. For example, for a threshold temperature of 37°C a reference voltage of 370 mV (0.37V) is needed. The full circuit diagram for the temperature alarm is shown in Figure 2.

IC1 is the temperature sensor and this is very easy to use. It is a basic three terminal device (two supply leads plus the output) and it will operate over a wide supply voltage range of 4 to 20V. Its current consumption is a mere 56 microamps, which together with a 5V supply potential gives totally insignificant self-heating.

IC2 is an operational amplifier which is used here as a voltage comparator. VR1 provides a reference voltage that can be set anywhere from 0V to approximately 1V, which matches the output voltage range of IC1. This reference voltage is applied to the inverting input of IC2 and the output of IC1 is coupled to the non-inverting input. Consequently, the output of

IC2 is low if the output of IC1 is below the reference voltage, or high if the output of IC1 exceeds the reference voltage.

When the output of IC1 is very close to the reference voltage there is a risk of instability, with a lot of noise being produced at the output of IC2. There are several factors that can contribute to this problem, but the main cause is the noise generated within IC2 itself. The positive feedback provided by R4 results in the output of IC2 switching cleanly and rapidly once the output from IC1 reaches a high enough voltage. Under standby conditions, with the output of IC2 low, R4 pulls the voltage at the non-inverting input of IC2 slightly lower. As the output switches to the high state, R4 tends to pull the voltage at IC2's non-inverting input fractionally higher. This produces a slight difference between the input voltage at which IC2 triggers to the high state and the input voltage at which it switches back to the low state again. This slight reluctance to change states (known as 'hysteresis') is sufficient to prevent instability.

The low frequency oscillator is based on IC3, and this is a standard 555 astable circuit. It is gated via the reset input at pin 4, which holds the output at pin 3 low when IC3 is gated off (when the output of IC2 is low). This prevents IC4 from oscillating. This is another 555 astable circuit, again gated via its reset input. It has an operating frequency of approximately 2.5kHz. When IC3 is activated, its output provides a roughly squarewave signal at a little over 1Hz. This pulses IC4 on and off, resulting in LS1 being fed with bursts of tone from IC4. This produces the alarm signal from LS1, which is a ceramic resonator.

Construction

Normally, a PC expansion card has to be a double-sided PCB. However, in this case the circuit is reasonably simple and it only requires connections to the ground and +5V terminals of the expansion bus. A single-sided board is therefore sufficient for this project. The printed circuit component overlay appears in Figure 3.

The CA3130E specified for IC2 is a MOS device, and it therefore requires the standard anti-static handling precautions. In particular, a holder should be used for this component, but it should not be fitted into the holder until the board is finished in other respects. Until then, it should be left in its anti-static packing.

I would not recommend the use of anything other than a CA3130E for IC2. Most other operational amplifiers will not work properly on a supply potential as low as 5V and many of those that will work do not provide a low enough output voltage to hold IC3 in the reset state. Some devices that will otherwise operate properly on a 5V supply tend to suffer from 'latch-up' problems when used in a comparator circuit (including the

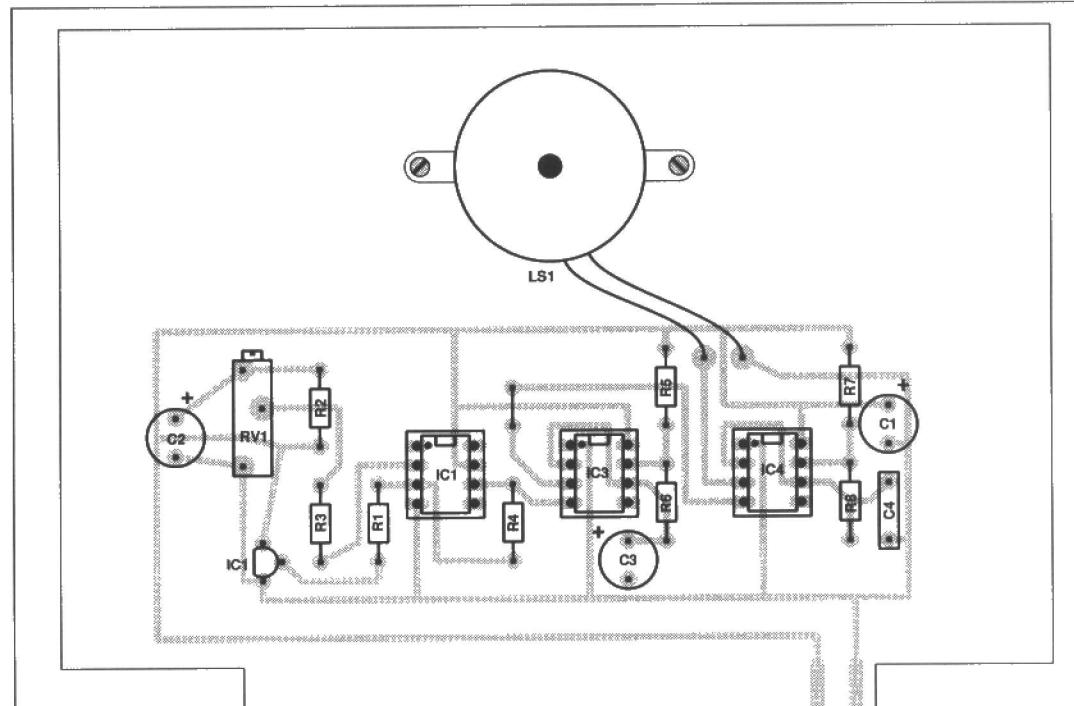


FIG.2. The P.C.B. component overlay

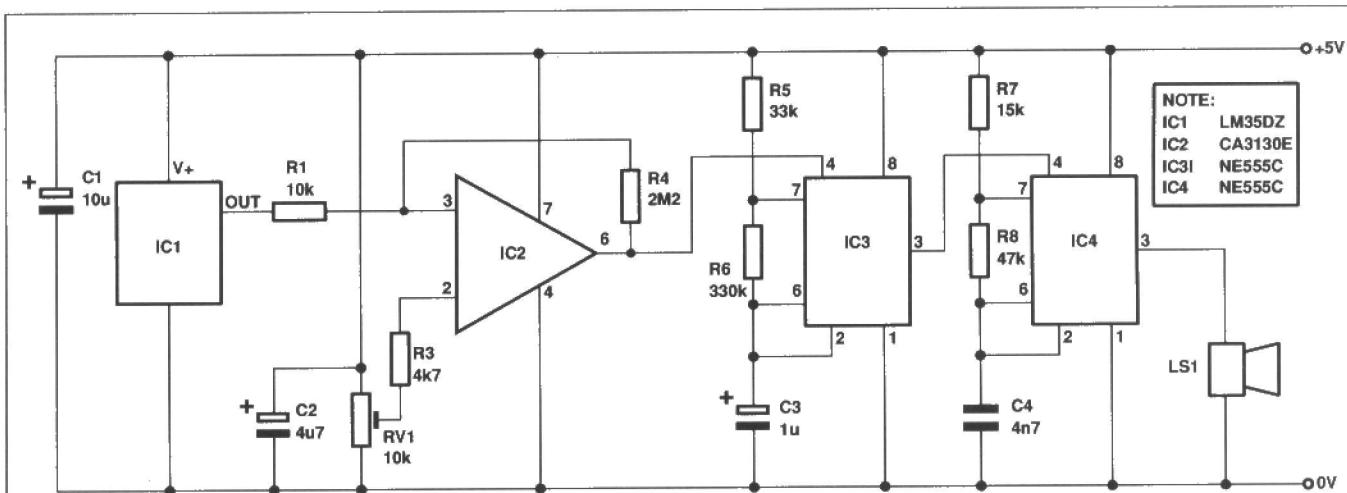


FIG.3. The PC temperature alarm circuit diagram

CA3140E). It is unlikely that satisfactory results will be obtained using anything other than a CA3130E.

Construction of the board is largely straightforward, but there are a few minor points that are worthy of note. Do not overlook the single link-wire just above R4. In order to fit into this layout properly, VR1 must be an 18 turn Cermet preset. C4 should be a type having 7.5mm (0.3in) lead spacing. LS1 is a cased ceramic resonator. Connections to this component are made via its 'flying' leads. Although the leads are often red and black, this is not a polarised component and the leads can be wired to the board either way round. It is mounted on the board via two 8BA (or metric M2) screws and mounting nuts. These are not normally supplied with the resonator. Alternatively, it can simply be glued in place using any general purpose adhesive.

I would not recommend the use of an ordinary moving coil loudspeaker with this project. Performance could be a bit erratic and even with a high impedance loudspeaker, quite high output currents could flow from IC4.

Adjustment

The alarm can be fitted into any spare expansion slot of the PC, but be careful to fit it the right way round. C2 must be towards the front of the computer, C1 and C4 must be towards the rear.

Before setting VR1 for a suitable threshold temperature, you must obviously decide just what that temperature should be. The technical specification in your computer's manual might provide some assistance here. Practical tests on a few PCs tend to suggest that their internal temperature is normally about 6 to 7 degrees above room temperature. For PCs having one of the more advanced processors such as a 80486, the difference is perhaps a couple of degrees more than this.

If we assume that the room temperature will not normally be much more than about 25°C, the interior temperature of the computer would normally be no more than about 35°C. Unless you have good reason to use a different threshold temperature, VR1 should therefore be set for a wiper potential of 350mV.

Trial and error can be used in the absence of test equipment to enable VR1 to be set for a suitable reference voltage. I suppose that this is actually the more reliable method of setting VR1, but it could be a bit time consuming. It is just a matter of finding the setting for VR1 that gives the lowest wiper voltage without the alarm sounding once the computer has been operating for a while (i.e. set the wiper of VR1 as far down the track as possible without the alarm activating once the computer has warmed up).

There is a slight complication in that the computer's outer

casing must be at least partially removed to provide access to VR1, but once VR1 has been adjusted, the outer casing must be put back into place so that the interior of the computer can warm up in the normal way. You must, therefore, allow time for the temperature inside the computer to rise back to its normal operating level each time VR1 is readjusted.

With a little experimentation, it should be possible to find a setting that does not give problems with false alarms, but still sets a low enough threshold temperature to trigger the unit if the computer's interior temperature becomes significantly higher than normal.

PARTS LIST

Resistors

(0.25 watt 5% carbon film)

| | |
|-----|---------------------|
| R1 | 10K |
| R2 | 39K |
| R3 | 4k7 |
| R4 | 2M2 |
| R5 | 33k |
| R6 | 330K |
| R7 | 15k |
| R8 | 47K |
| VR1 | 10k 18 turn trimpot |

Capacitors

| | |
|----|---------------------------|
| C1 | 10 μ 25V radial elect |
| C2 | 4u7 50V radial elect |
| C3 | 1 μ 50V radial elect |
| C4 | 4n7 polyester |

Semiconductors

| | |
|----------|---------|
| IC1 | LM35DZ |
| IC2 | CA3140E |
| IC3, IC4 | NE555C |

Miscellaneous

LS1 Cased ceramic resonator
Printed circuit board
8 pin DIL IC holder (3 off)
Wire, solder, etc.

OMNI ELECTRONICS

174 Dalkeith Road, Edinburgh EH16 5DX • 031 667 2611

The supplier to use if you're looking for —
★ A WIDE RANGE OF COMPONENTS AIMED
AT THE HOBBYIST ★
★ COMPETITIVE VAT INCLUSIVE PRICES ★
★ MAIL ORDER - generally by
RETURN OF POST ★
★ FRIENDLY SERVICE ★



Open: Monday-Thursday 9.15-6.00
Friday 9.15-5.00 Saturday 9.30-5.00



Hesing Technology



41, Bushmead Road, Eaton Socon
Huntingdon, Cambs. PE19 3BT

Tel: (0480) 386156
Fax: (0480) 386157

TEST EQUIPMENT

Supply
Maintenance
Commissioning

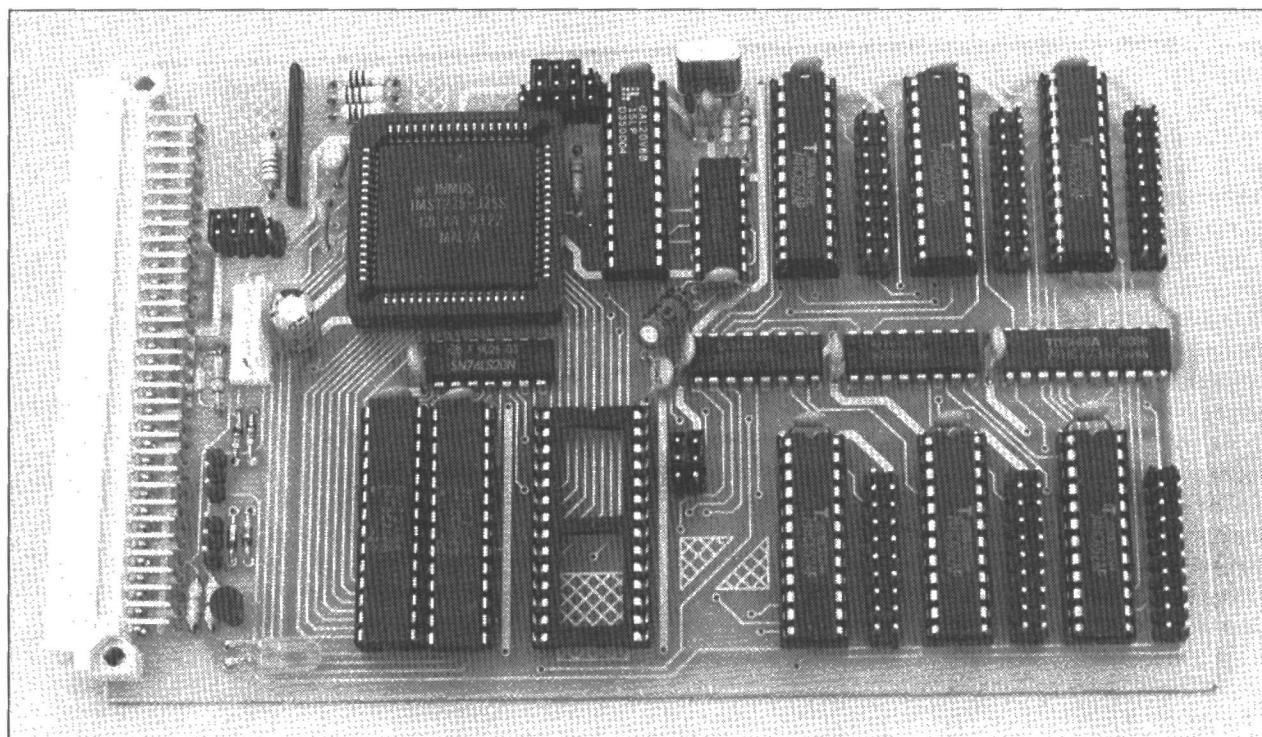
SYSTEM CONSULTANCY

Replacement Parts
Supply of Service &
Operators Manuals
Components

Distributors for:
WAUGH INSTRUMENTS, RAMTEST LTD, KRENZ ELECTRONICS, PANTHER

Transputer Based Si

- An Introduction



Fancy building your own supercomputer or sophisticated graphics engine? Then why not try this versatile single board design from Andy Papageorgiou and Mark Robinson

Back in September 1985, Inmos unveiled the transputer amid a blaze of publicity, promises of unlimited computer power and cute ray-traced Newton's Cradles on Tomorrow's World. Nine years later and it seems that the transputer has been mostly forgotten - headlines today are dominated by the new Intel processors. Nevertheless, transputers found a niche in highly parallel supercomputers (the Meiko computing surface, for example), graphics processors and research machines, which justified their continued development by Inmos (now owned by SGS-Thomson). The imminent release of the 64-bit T9000 proves that the transputer is far from being consigned to computing history. The current range of transputers is described in Table 1.

The Von Neumann Machine

Long before the first electronic computer was ever built, mathematicians were concerned with the problem of computability, the question of whether a given problem could be solved in a

finite time. The quest to answer this question led to the development of then hypothetical computing machines, such as the Turing Machine and the Von Neumann Machine. Eventually, electronics caught up with the mathematicians' desires and computing machines began to appear, based on the Von Neumann architecture as it offered the most practical design.

A block diagram of the Von Neumann machine is shown in Figure 1 (with modern labels). It is an example of a 'Finite State Machine' - the CPU can only be in one of a finite number of states, and its action depends both on the current input and the current state. The program counter is central to the Von Neumann machine's operation, it points to the position in memory of the current instruction. Also, the data input and output share the same memory, which allows the current action to depend on outputs generated a long time in the past. It can

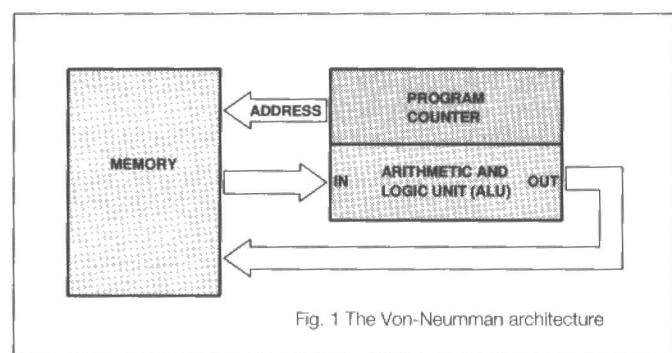


Fig. 1 The Von-Neumann architecture

Single Board Computer

be proved that any computable problem can be solved on a Von Neumann machine (given an infinite amount of memory).

As memory devices became larger and faster and program sizes grew from a handful to many millions of instructions, it soon became apparent that the Von Neumann architecture had problems. Von Neumann CPUs can only process one instruction at a time, all the other instructions just sit waiting in memory, in the same way that cars on the M6 sit waiting when two lanes have been coned off to park spare JCBs in. Such traffic jams are referred to as bottlenecks by the helpful traffic reports on the radio and naturally the CPU's serial nature became known as the Von Neumann bottleneck. Much research effort has been put into increasing computers' throughput and the most successful approaches are summarised in Figure 2. The first scheme is to widen the data bus, increasing the size of number that can be dealt with by a single instruction (Figure 2b). The second approach is to increase the speed at which instructions can be processed, either by increasing the clock speed of the chip, or using clever hardware tricks like pipelining, vector processing and RISC architectures (Figure 2c).

Obviously, these two methods are limited by the technology available and power CPUs today are fast approaching the limits of what can be achieved in silicon. The parallel computer (Figure 2d) is not limited by technology, only by the number of pro-

cessors that you can afford. Parallel computers will always have the edge over serial machines since, to coin a phrase, two heads are better than one. It is no surprise that today's supercomputers are parallel or even massively-parallel machines. It has been estimated that the performance of a parallel machine is five years ahead of that of a serial machine.

Parallel Computing and the Transputer

A parallel computer is not simply a large number of CPUs. If it were, then ten PCs in the same room would be a parallel computer. The key to successful multiprocessing is communication - connect those ten PCs to a network and tasks can be distributed amongst them, exchanging data whenever it is required. This sort of multiprocessor, where each processor is a complete computer in its own right, is normally called a distributed system, but the principle is the same. Without efficient communications channels there can be no parallel processor.

Early attempts to make parallel computers failed because inter-processor communication needed to be implemented in hardware, usually via some kind of shared memory connected to each processor by tristate bus transceivers. The extra hardware costs involved, together with the difficulty in routing the masses of databases required around the PCB and the need to implement software to ensure that the right data was in the right place at the right time, meant that parallel machines were not commercially viable.

This is where the transputer comes in. On each transputer chip is a basic CPU core, a small amount of RAM and, most importantly, a number of high speed serial communications links, capable of transferring data between transputers at 20Mbits/sec over two wires. Most transputers have four links, which allow large, two dimensional arrays of processors to be constructed, each capable of communicating with its four nearest neighbours (Figure 3). To support the links, the transputer instruction set contains a number of instructions to simplify the sending and receiving of data on the links.

Processes and Scheduling

Any new technology brings with it a new set of jargon words and transputers are no exception. The first important technical term is a

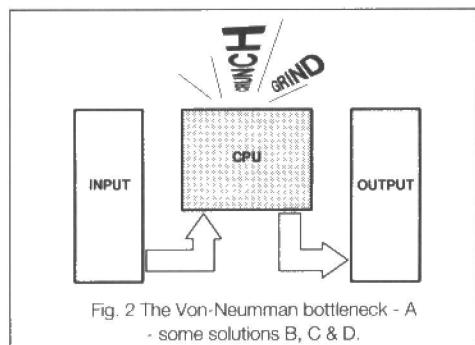
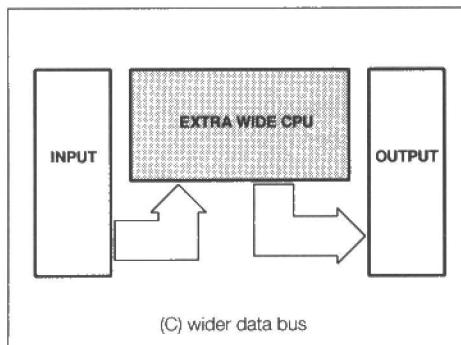
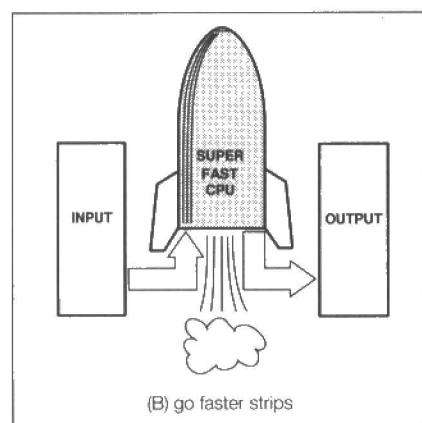


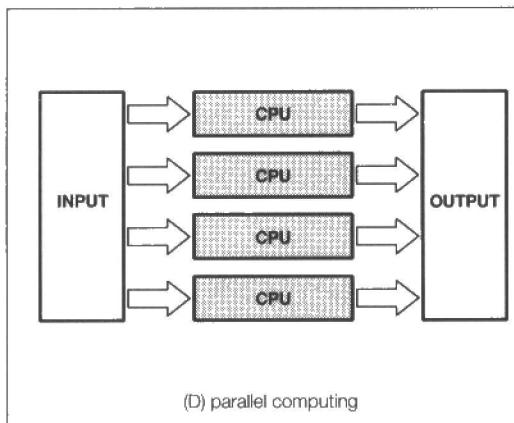
Fig. 2 The Von-Neumann bottleneck - A - some solutions B, C & D.



(C) wider data bus



(B) go faster strips



(D) parallel computing

'process' which is simply a fragment of code which performs some action, like a subroutine in a high level language. A process may be a collection of smaller, simpler processes and the complete set of processes forms the program. Whereas subroutines are called one after the other by the main program, a parallel computer may have several processes running at the same time, or concurrently in the jargon, on several different processors.

Scheduling is the job of deciding the running order (or schedule) of the processes that make up a program. The idea is to spread the processes out over the available processors in such a way as to minimise the time processors spend idling while processes wait for others to provide data. The best way to schedule processes is dependent upon the number and connectivity of processors, two processes which need to communicate must run on processors which have a link directly connected between them. Large arrays of transputers often have the links routed via a switching matrix, to allow the link connections to be configured by software to suit the task in hand.

The task of optimally scheduling a set of processes onto a given transputer array is NP-complete, which means that the only way to be sure that you have found the best schedule is to test every possible schedule, which is impractical for anything other than tiny programs on a few processors. Fortunately, we can use rules of thumb to choose a good (although not necessarily the best) schedule. The problem of optimal scheduling is currently a hot research topic, keeping parallel architecture groups busy in universities the world over.

Consider the example of adding a large column of numbers. There are many ways that this task could be scheduled, but common sense tells us that a split-and-merge technique (Figure 4) is the best. Each processor is given a share of the numbers to sum, and when they have finished the subtotals are collected and added to provide the final answer. This simple example shows how intertwined the scheduling is with the hardware configuration. The optimum schedule is that which ensures all the subtotals become available at the same time, so the more powerful processors should get a bigger share. The interconnections become important when summing the subtotals. In the figure, each processor can only talk to its nearest neighbour, so the best approach is for CPU2 to accept CPU1's subtotal, add it's own and passes the result to CPU3 and so on. It turns out that, for adding a column of numbers, Figure 4 is the best possible arrangement.

Multitasking

There was a time when multitasking was confined to large mainframes and workstations running UNIX. Today most people are familiar with the concepts of multitasking, largely due to the popularity of Microsoft Windows, a multitasking operating

system for PCs. Multitasking is the ability of a single processor to apparently run several processes concurrently. What actually happens is that the processor runs each process individually for a short time before shelving it (called descheduling, not to be confused with scheduling a program as described above) and moving on to the next. Obviously, when a process is descheduled (a 'context switch' in the jargon), its current state (the 'context') must be saved somewhere to allow it to be picked up from where it was left off when it is next scheduled.

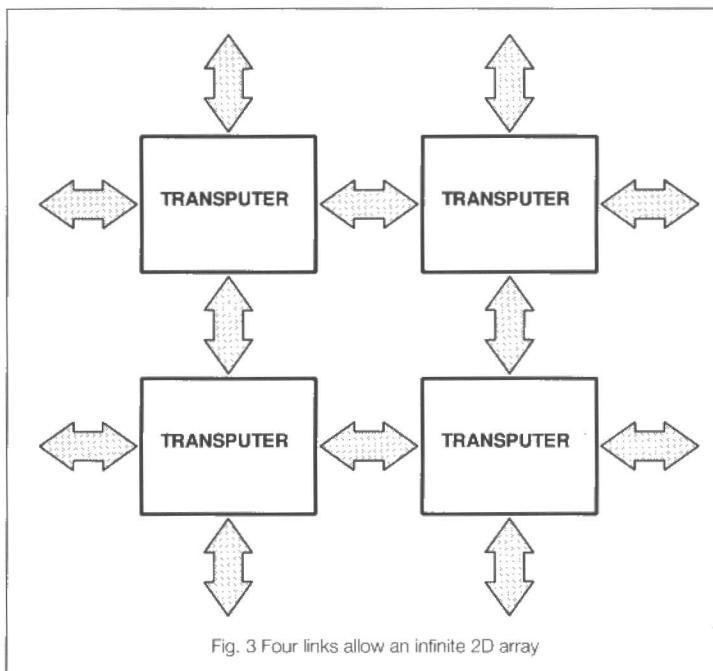
Any microprocessor can be programmed to multitask, however the programmer must write a complicated scheduler to properly allocate time slices to the processes which want to run, to ensure that processes don't conflict for memory and to handle context switching. Context switching usually involves saving the register set and any workspace in a temporary store and loading the previously saved values for the new process, which can take a considerable time, reducing the performance of the system. The programmer must also be careful to ensure that his processes are written in such a way that they give control back to the scheduler when their time is up.

On the transputer, multitasking is handled in hardware with very little overhead. For example, a 25MHz T225 can achieve a context switch in less than 800ns. The instruc-

tion set contains instructions to simplify the creation and control of concurrent processes. Scheduling of processes is carried out automatically and transparently, a process is descheduled whenever it cannot continue (it may be waiting for data from another process for example). When there is more than one active process, these are automatically timesliced.

The transputer allows two levels of process priority. A normal, low priority process is one which can run concurrently, while a high priority process takes over the CPU entirely and runs until it terminates, or until it requires input from another process. High priority processes are designed for short, time critical operations.

A novel design feature of the transputer is that, as far as the programmer is concerned, the method of communication between processes is the same, regardless of whether they are running concurrently on one transputer or on separate transputers and communicating over a link. Communication takes place over a 'channel', a pre-arranged pair of memory locations through which data is transferred. For communication between two processes on the same transputer, the address of these memory locations (the channel address) is chosen by the programmer. For communication over links, one of eight special reserved addresses in the bottom of the memory map is used. This means that code can be developed on one (or a few) transputers and distributed around a large array later with very little modification.



The Register Set and Workspace

The transputer CPU contains six registers, labelled A, B, C, W (workspace pointer), I (instruction pointer) and O (operand register). The function of the operand register will be described in conjunction with the instruction set in the next section. The I register is the program counter, that is it points to the next instruction to be executed. The A, B and C registers form a three level push down stack used for all arithmetic and logical operations. The A register is the top of the stack and the C register the bottom and all arithmetic and logical operations implicitly refer to the stack. For example, the 'add' instruction pops the top two values off the stack, adds them and pushes the result onto the stack. Programmers of reverse Polish languages (e.g. FORTH or PostScript) will be used to using an evaluation stack, the rest of us have to adapt from our accumulator and registers view of programming.

The W register holds the address of the currently active process's workspace. Every process has its own workspace which is used to store local variables. The instruction set provides instructions to load and store relative to the workspace pointer. The first 16 words above the workspace pointer can be accessed in a single cycle and are effectively equivalent to 16 general purpose registers. Any memory location can be referenced relative to the workspace pointer, allowing workspaces to be any size. It is the programmer's responsibility to ensure that processes don't overwrite each other's workspaces.

It is this workspace arrangement that allows the transputer to context switch so quickly. To perform a context switch, the transputer only needs to load the W register with the address of the new workspace and the I register with the address of the instruction that was about to be executed when the process was last descheduled. This address is saved in a reserved part of the workspace when a process is descheduled.

When a process is rescheduled, the contents of the stack and O register will be undefined. This doesn't matter though, because rescheduling is only allowed after certain instructions which guarantee that these registers don't contain any valid data.

The Instruction Set

Large arrays of transputers are usually programmed in a high level language, in particular Occam, the language developed by Inmos specifically for parallel programming on transputers. While Occam simplifies complicated program development, small systems built around T2 series transputers where memory capacity is limited are best programmed in the transputer's native assembly language. The instruction set is somewhat unconventional and a brief look would be interesting and instructive. What follows is not intended as a complete guide to transputer assembly language, for that consult the Inmos datasheets or the books recommended later. On all the present transputers, instructions are eight bits wide, split into 4 function bits and a four data bits. This allows for 16 functions on numbers between 0 and 15. Thirteen of these functions are reserved for the basic computing operations, called the direct functions, listed in Table 2.

This may seem a little limiting, but the remaining three functions are used with the O register to extend both the number of functions and the size of number that can be handled.

All direct functions place their data bits into the lower four bits of the O register, which is then used as the operand for the instruction. The Prefix instruction also loads its data into the O register, but then it shifts it up four bits (the Negative Prefix instruction complements the O register first). This allows large operands to be built up, for example the instruction

LDC \$4AF3 ; Load 4AF3H to A register would assemble as

PFIX \$4 PFIX \$A PFIX
\$F LDC \$3

The remaining function code, Operate, causes the transputer to decode the O register as if it were an instruction and execute it. These instructions,

called indirect functions, all operate on the stack and hence don't need an operand. For example, the instruction

DIV ; Divide - A reg :=
A reg / B reg

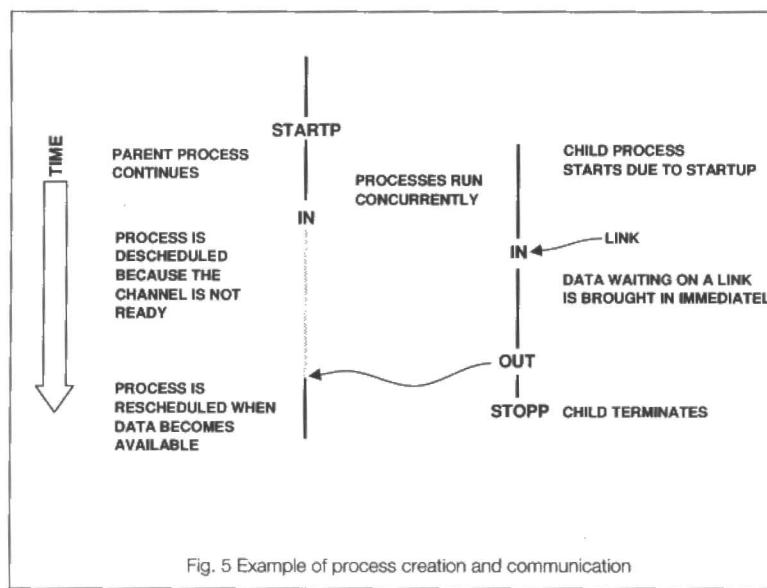
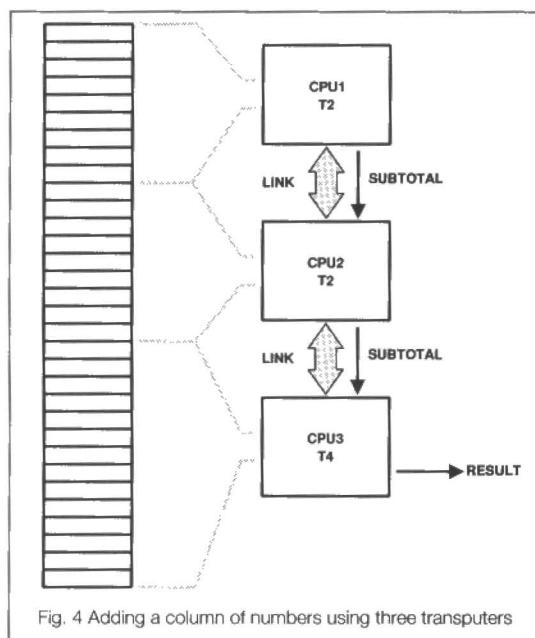
which has the opcode 2CH, would be encoded as

PFIX \$2 OPR
\$C

The PFIX and OPR instructions are normally added automatically by the assembler and are transparent to the programmer.

Inmos claims that this instruction format leads to very efficient code, since 70% of a typical program will consist of single byte and hence single cycle, instructions.

Creating Processes



Process control is performed using the STARTP, STOPP and RUNP instructions. A process is created (forked) using the following code

```

LDC newproc - here      ; Offset to
start of new process
LDC wksp    ; Address of new process
workspace
STARTP      ; Fork the new process
here: ..... ; parent process continues
newproc: ..... ; child process start

```

After execution of the STARTP instruction, there will be two processes running concurrently. The code after the 'here' label will run as the parent process, while the code after the 'newproc' label will run in parallel as a second child process.

Communication between processes is just as easy. To receive data on a channel, the following code is used

```

LDLP 50      ; store message 50
bytes above workspace
LDC channel ; channel address
LDC bytes   ; number of bytes in
the message
IN      ; get message

```

If the message is not ready immediately, the process will be descheduled until the channel becomes active. The channel address can be any memory address, unless communication over a link is required when one of the special channels must be used. The transmitting process uses the same code except that the OUT instruction is used rather than the IN instruction. The effects of these instructions are illustrated in Figure 5.

Conclusions

The transputer is a unique microprocessor and hopefully this short article has highlighted the ease with which multiprocessor or multitasking systems can be designed using it. The transputer is now a mature technology and prices are well within the reach of the hobbyist. Next month, ETI describes a constructional project to build a single board microcontroller using a T225, believed to be the first project to use a transputer published in a hobbyist magazine.

The card forms a useful controller in applications where 8-bit micros aren't powerful enough, image processing or neural networks for example, as well as being a good evaluation board to experiment with transputers.

For further information on the transputer, Inmos produces a range of datasheets and applications books: Inmos, Planar House, Parkway Globe Park, Marlow, Bucks, 0628 890800. In particular The Transputer Instruction Set: A Compiler Writer's Guide is useful for programmers of assembly language.

The Transputer Handbook by Ian Graham and Tim King (available from Maplin) may be worth a look. It contains details of the whole transputer family, the assembly language and Occam.

Readers with access to the Internet may care to browse around the parallel archive at Lancaster University. FTP to unix.hensa.ac.uk and look in the directory '/parallel'. Also, lively discussion about transputers, Occam and parallel programming can be found on the USENET newsgroup comp.sys.transputer.

The complete transputer family

| Processor | BusWidth | RAM | Link |
|-----------|----------|-----|-------------------------------------|
| Comments | | | |
| T212 | 16-bit | 2K | 4 |
| | | | Obsolete - use T225 |
| M212 | 16-bit | 2K | 2 |
| | | | T212 with internal disk controller |
| T222 | 16-bit | 4K | 4 |
| | | | Obsolete - use T225 |
| T225 | 16-bit | 4K | 4 |
| | | | Entry level 16 bit transputer |
| T400 | 32-bit | 2K | 2 |
| | | | Low cost 32-bit transputer |
| T414 | 32-bit | 2K | 4 |
| | | | Obsolete - use T425 |
| T425 | 32-bit | 4K | 4 |
| | | | 32-bit Integer transputer |
| T800 | 32-bit | 4K | 4 |
| | | | Obsolete - use T805 |
| T801 | 32-bit | 4K | 4 |
| | | | 32-bit + 1MFLOP FPU, SRAM interface |
| T805 | 32-bit | 4K | 4 |
| | | | 32-bit + 1MFLOP FPU, DRAM interface |

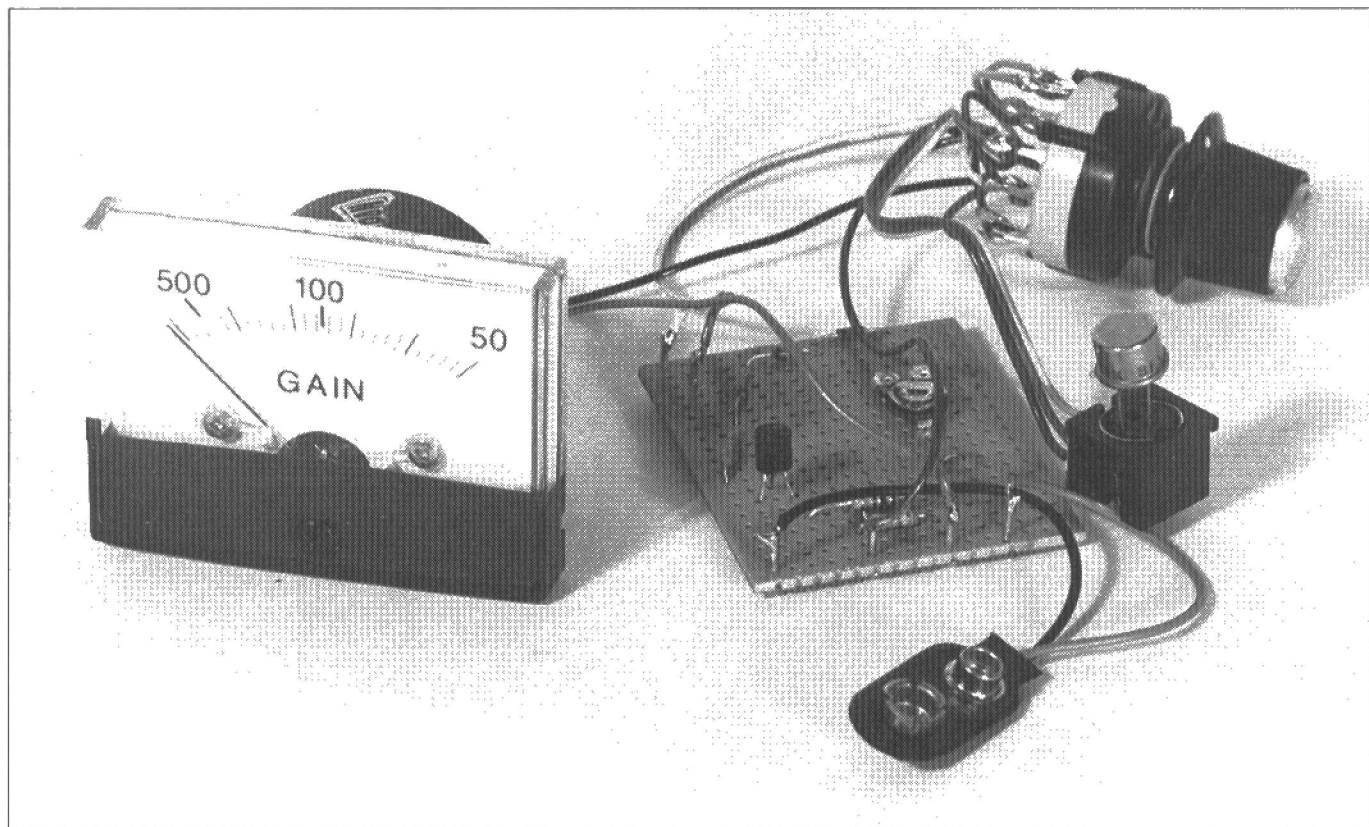
The 16 basic functions

| Opcode | Name | Function |
|---------|------------------------|---|
| LDC N | Load Constant | Push N onto stack |
| LDL N | Load Local | Push word stored at the address N words above the workspace to the stack |
| STL N | Store Local | Pop A and store at the address N words above the workspace |
| LDNL N | Load Non-local | Push word stored at the address N words above the base address to the stack |
| STNL N | Store Non-local | Pop A and store at the address N words above the base address |
| LDLP N | Load Local Pointer | Load the address which is N words above the workspace |
| LDNLP N | Load Non-local Pointer | Load the address which is N words above base memory |
| ADC N | Add Constant | $A := A + N$ |
| EQC N | Equals Constant | If $(A = N)$ then $A := 1$ else $A := 0$ |
| J N | Jump | Relative Jump |
| CJ N | Conditional Jump | If $(A = 0)$ then take Relative Jump (A is popped) |
| AJW N | Adjust Workspace | Move Workspace pointer by N words |
| CALL | Call | Relative Subroutine Call |
| PFIX | Prefix | |
| NFIX | Negative Prefix | |
| OPR | Operate | |

Next month...

We will start on the practical aspect of actually building the transputer board.

CONSTANT CURRENT TRANSISTOR TESTER



Make transistor testing easy with this useful little device from Robert Penfold

The vast majority of transistor testers operate by feeding a fixed base current into the test device and then measuring the resultant collector current. For example, using a base current of one microamp and a 1mA meter in the collector circuit, the current gain of the transistor is equal to the indicated collector current in microamps. This method has the advantage of providing a linear and forward reading scale, so there is no need for any alterations to the meter's scale. It is a method that suffers from a major drawback however, and this is simply that the gain measurement is not made at a particular collector current. If you test an assortment of ten different transistors,

their gains will probably be measured at ten significantly different collector currents. The higher the gain of the test device, the higher the collector current that flows while the measurement is made.

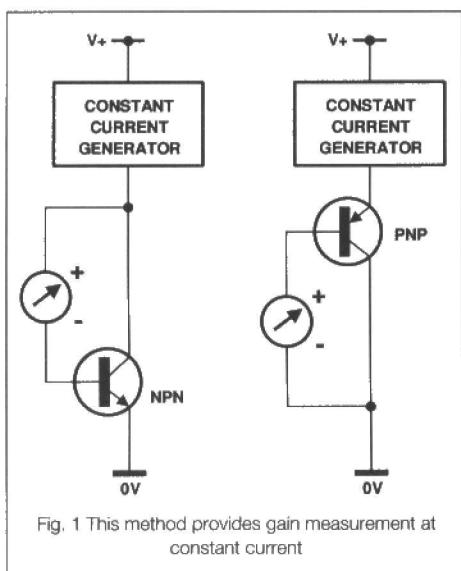
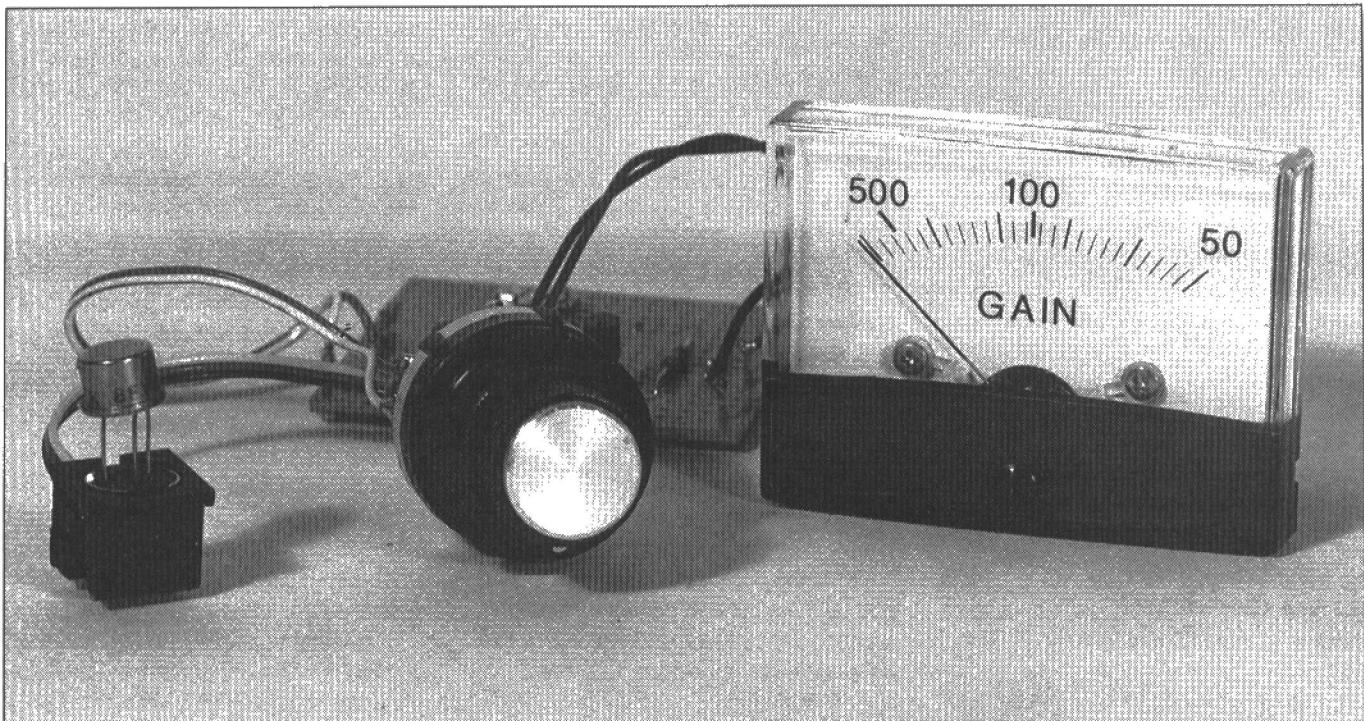
This factor is not purely of academic importance. The gain of a transistor changes significantly with variations in its collector current and in general, the higher the collector current, the higher the current gain. In this case, it is the fact that small collector currents produce very low current gains that is of significance. When testing something like a high gain audio transistor, a relatively high collector current will flow. The exact current obviously depends on the design of the tester and the current gain of the test device, but it would typically be between about 0.5 and 5mA. This is high enough to ensure that the test device operates efficiently.

Problems tend to arise when testing something like fairly low gain switching or

r.f. transistor. Devices such as these have gains that are usually around 50 to 100 or so, and they are mainly intended for operation at collector currents of several millamps. Due to their relatively low gains, when checked using a conventional transistor tester they are likely to operate at a collector current that is only a fraction of a millamp. The low collector current gives reduced current gain, which in turn results in an even lower collector current. This vicious circle results in certain types of transistor being measured at very low collector currents, which produces unrealistically low gain readings. The user can allow for this factor when assessing results, but when testing low gain, medium current devices, most conventional transistor testers provide ambiguous results.

How It Works

There is a simple alternative which ensures that all serviceable test devices are measured at the same collector



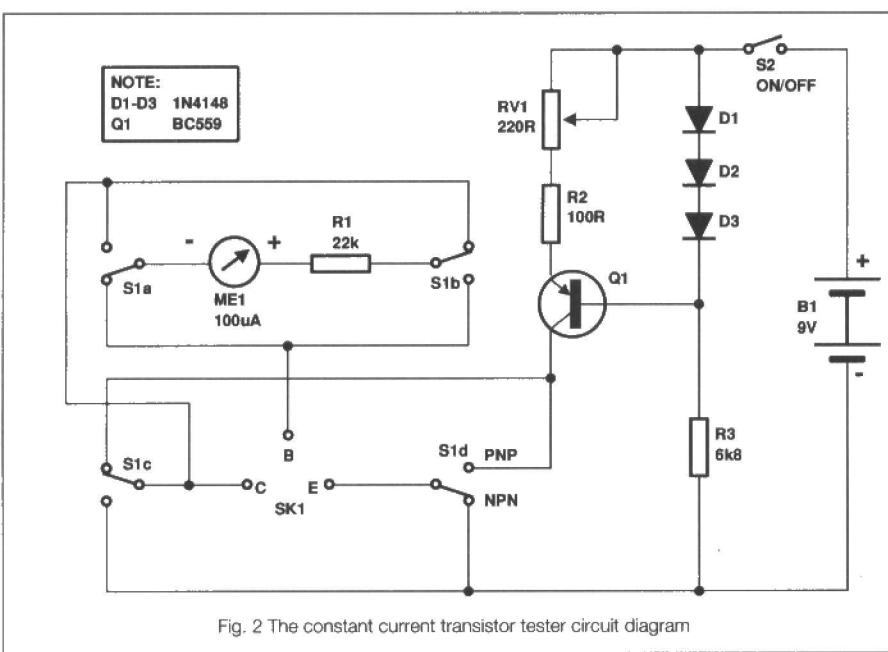
current. The basic arrangement used with this system is outlined in Figure 1. The left hand diagram shows the set-up used for testing npn transistors. The pnp version is shown on the right hand side and is essentially the same, the only difference being that the collector and emitter terminals have been swapped over, so the test transistor is connected to the supply with the correct polarity. This necessitates a switch in the polarity of the meter as well.

The test device has its collector and emitter terminals fed from the supply rails via a constant current generator. A meter is connected between the collector and base terminals of the test component and due to the low resis-

tance through the meter, it provides a strong base current to the test device. This results in the test component being biased hard into conduction, but the constant current generator limits the current flow to a safe level. The important factor here is that the base current which flows into the test device will be equal to the collector current divided by the gain of the transistor. The higher the gain of the transistor, the lower the base current. Instead of providing a constant base current and measuring the collector current, the test component is fed with a constant collector current and it is the resultant base current that is measured.

This arrangement measures the gain of every test component at virtually the same collector current. Slight variations occur because some of the current flows into the base circuit of the test device and the amount of current that is tapped off by the base circuit depends on the current gain of the transistor. However, this only produces variations of about one or two percent, which is not enough to significantly affect the accuracy of the system.

The main drawback of this method is that the meter has a non-linear and reverse reading scale, rather like an analogue multimeter when used on a resistance range. Suppose that measurements are made at a collector current of 5mA, and that the meter has a full scale value of 100 μ A. Full scale deflection of the meter corresponds to a current gain of 50 (5000μ A/ 100μ A = 50). At half full scale deflection, the base current is only half as much, so the current gain must



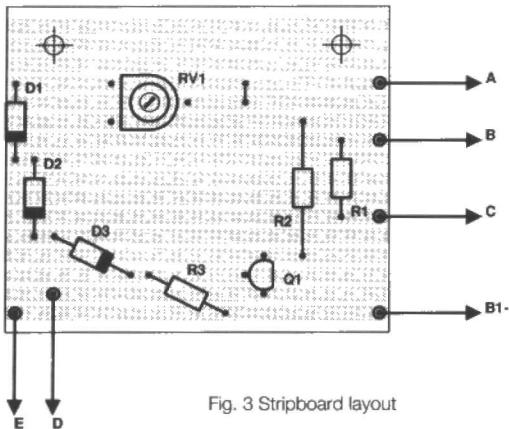


Fig. 3 Stripboard layout

be twice as high ($5000\mu\text{A}/50\mu\text{A} = 100$). A meter reading of $10\mu\text{A}$ corresponds to a current gain of 500 ($5000\mu\text{A}/10\mu\text{A} = 500$). The scale is well spread out at low current gains, but is rather cramped towards the high gain end. Provided that a sensible collector current and meter sensitivity are chosen, as in this example, results are still perfectly usable, however.

Figure 2 shows the circuit diagram for the constant current transistor tester. The constant current generator is based on Q1, and uses a conventional arrangement. R3 plus D1 to D3 provide a stabilised potential to the base of Q1, which is about 1.8V below the positive

supply voltage. This gives about 1.2V across VR1 and R2. VR1 is adjusted to give a collector current of 5mA with the collector of Q1 short circuited to the earth rail. The output current remains virtually the same with higher collector load resistances, up to the point where Q1 can no longer provide a high enough collector voltage to produce a flow of 5mA. In normal operation, the load resistance is always low enough to provide the full 5mA collector current.

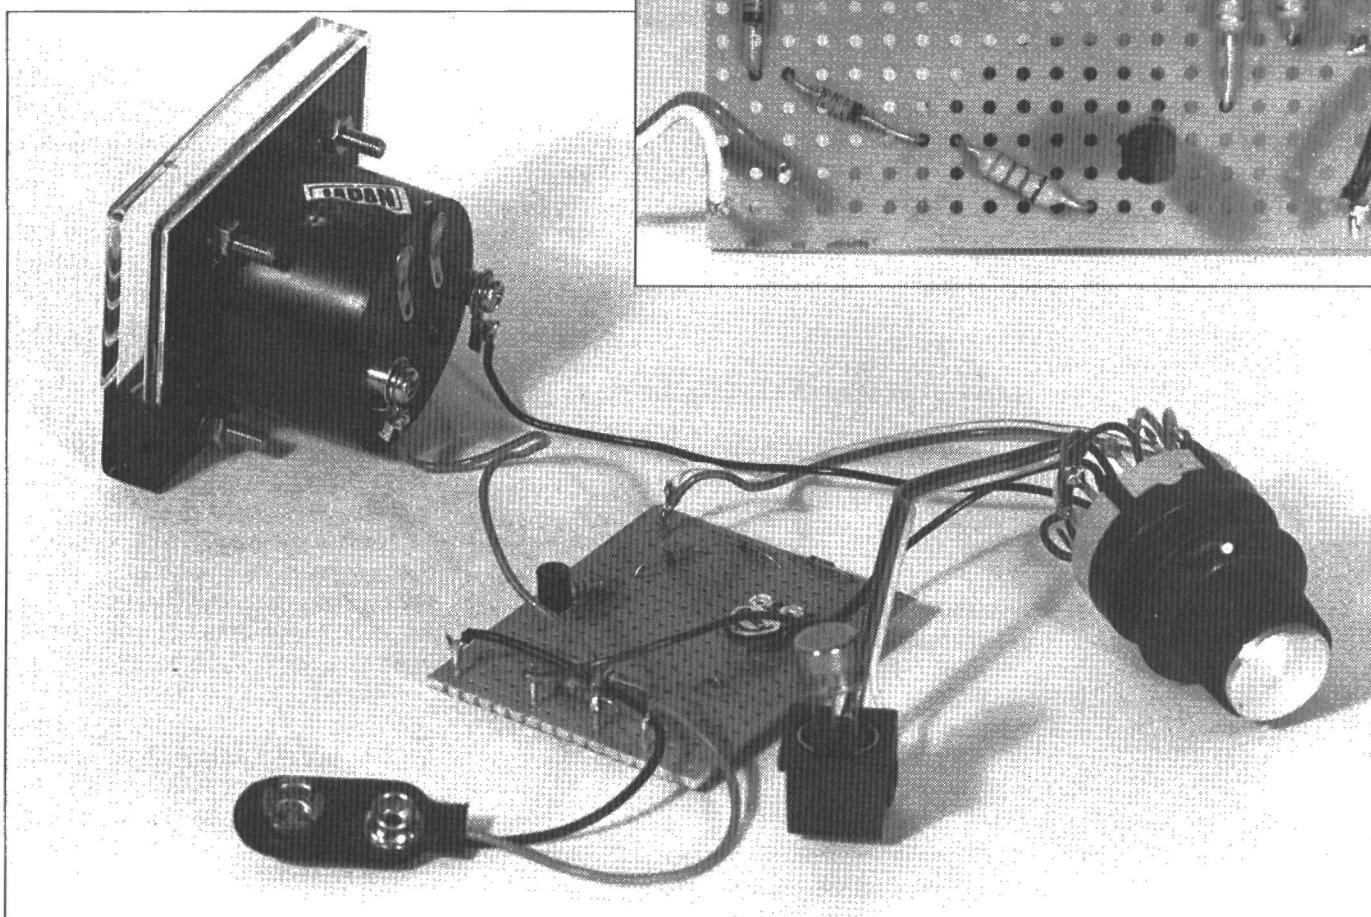
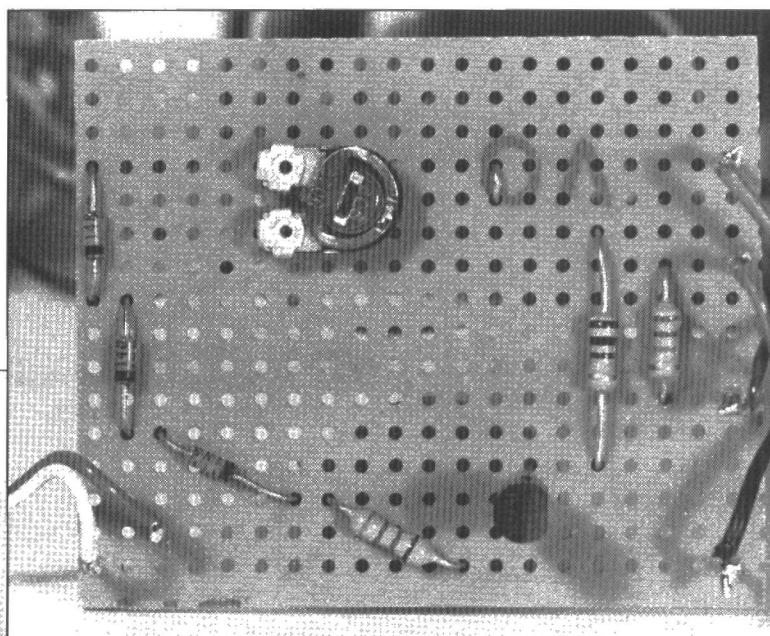
S1 is the npn/pnp mode switch. S1c and S1d are used to connect the test component's emitter and collector terminals to the current limited supply with the appropriate polarity. S1a and S1b ensure that the polarity of the meter is correct for the selected mode. If the test component should be faulty,

with a short circuit across the base and emitter terminals, the full 5mA output of Q1 could flow through ME1. R1 provides current limiting to ensure that ME1 cannot pass such a high current. It also raises the collector voltage at which the gain measurement is made. Without R1, the collector test voltage would be well under 1V.

The current consumption of the circuit is approximately 1mA under standby conditions. The consumption increases to about 6 mA when a transistor is connected to the test socket.

Construction

Figure 3 provides details of the component panel, which is based on a piece of stripboard measuring 20 holes by 16 copper strips. Construction of the board is very straightforward indeed. There are no breaks in the copper strips, but do



not overlook the short link-wire to the right of VR1.

The hard wiring is illustrated in Figure 4. This is rather more convoluted, but should not provide any serious difficulties provided that you set about things methodically. The main problem is the substantial amount of wiring to S1. This is a standard 3 way 4 pole rotary switch, but its adjustable end-stop is set for 2 way operation. I found it easiest to add the link-wires shown in the lower view of S1 and to then add the connections shown in the main wiring diagram. You might find it easier to do things the other way round, but either way it is necessary due take due care not to accidentally desolder any of the existing wiring when adding new connections.

I found that a 3 way miniature DIN socket was the most convenient type to use for SK1. Most small transistors will plug straight into this without any difficulty. A set of test leads will be needed in order to make the connections to any transistors which have unusual encapsulations and lead out arrangements (small power devices, for instance). It is advisable to use wires of three different colours so that the base, emitter and collector test leads are easily identified. I find it easiest to make the connections to the test devices via very small crocodile clips of the type which are largely PVC sleeved. Miniature probes of the spring-loaded variety are also suitable, but can be quite expensive.

The meter requires a large round mounting hole in the front panel. For a standard 60 by 48mm component the cut-out should have a diameter of 38mm. It can be cut using a fretsaw, coping saw, 'Abrafile', etc. The positions of the four smaller mounting holes can then be located using the meter itself as a sort of template. These four holes should be 3.2 mm in diameter.

Ideally, ME1 should be recalibrated so that it is directly calibrated in terms of current gain. The front of a modern panel meter simply unclips and it is usually necessary to do no more than remove a couple of small screws in order to free the scale plate. The

existing numbers can be carefully scraped off using a small modelling knife and then rub-on transfers used to add the new scale. This procedure is not particularly difficult, but the movement of a panel meter is quite delicate. If you decide to recalibrate the meter, take due care when the front cover is removed and the movement is exposed.

It is not essential to recalibrate the meter and the alternative is to remember the current gains at a few key points on the scale. This list of current values and corresponding current gains should be of help whichever method you select.

| CURRENT READING | Current | Gain |
|-----------------|---------|------|
| 100µA | 50 | |
| 70µA | 71 | |
| 50µA | 100 | |
| 40µA | 125 | |
| 30µA | 167 | |
| 25µA | 200 | |
| 20µA | 250 | |
| 15µA | 333 | |
| 10µA | 500 | |
| 5µA | 1000 | |

Adjustment and Use

VR1 must be adjusted to give an output current of 5mA from the constant current generator. Set S1 to the npn position (the counter-clockwise setting) and connect a multimeter across the collector and emitter terminals of SK1. The positive test prod connects to the collector terminal and the multimeter should be set to a DC current range having a full scale value of about 5 to 20mA. Switch on the transistor tester and adjust VR1 for a

reading of 5mA on the multimeter. The tester is then ready for use.

When connecting the unit to a test device via the set of test leads, the collector and emitter terminals should be connected first. A high current will flow through the meter if the base and emitter terminals are connected first, although this should not result in any damage to the meter. The problem does not arise when connecting test components direct to SK1, because all three leads are connected and disconnected (more or less) simultaneously.

| PARTS LIST | |
|-----------------------|--|
| Resistors | |
| R1 | 22k |
| R2 | 100R |
| R3 | 6k8 |
| VR1 | 220R min hor preset |
| Semiconductors | |
| Q1 | BC559 |
| D1,2,3 | 1N4148 (3 off) |
| Miscellaneous | |
| S1 | 3 way 4 pole rotary (set for 2 way operation) |
| S2 | SPST min toggle |
| B1 | 9V (PP3 size) |
| SK1 | 3 way Min DIN ME1 100µA moving coil panel meter, case, battery clip, 0.1in stripboard having 20 holes by 16 strips, test leads, control knob, wire, etc. |

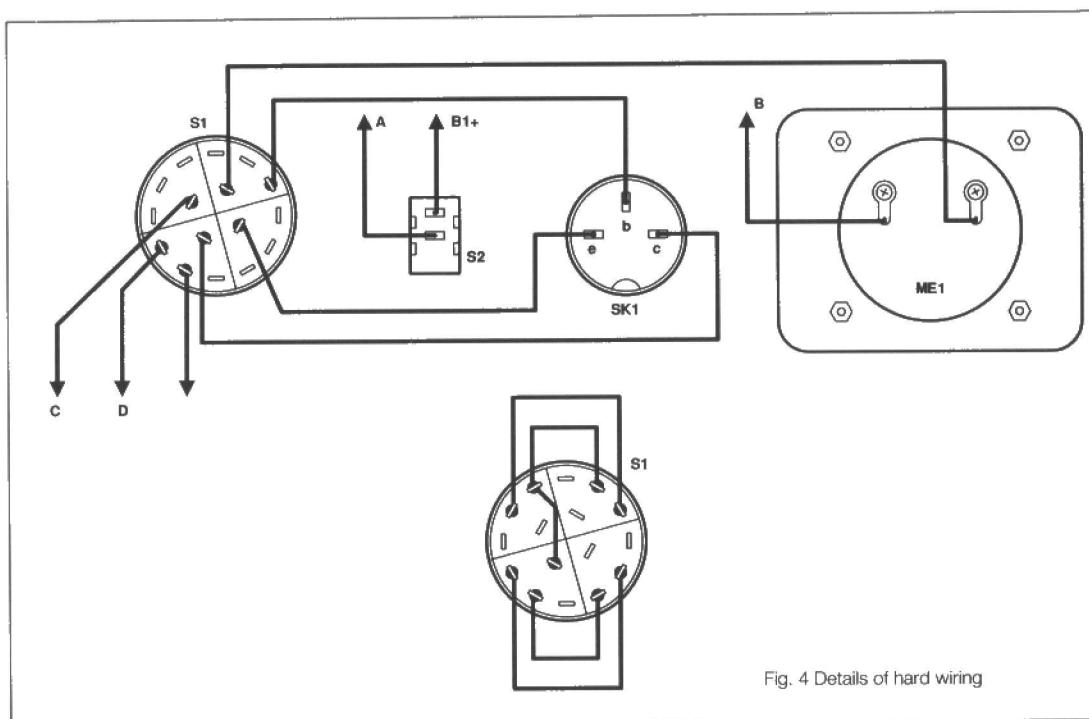


Fig. 4 Details of hard wiring

PC Clinic

This month in PC Clinic we examine the memory used in a PC, how it is configured, how it can be expanded and how faults can be located.

It is impossible to have a working computer system without some form of memory. Memory is used to store the sequence of commands that is processed by the CPU and it also stores the data which is used, or created, by the sequence of commands. How much memory is needed and what type of memory is used depends, however, upon the design of the system.

In theory, a wide range of different types of memory could be used. Given the existence of a few processor registers and some special control hardware, it would be possible to use a disk drive as the main processor memory. However, because access times for a disk drive are extremely slow in relation to any currently available processor, it would make the whole system slow since the processor would have to wait for every instruction and every piece of data.

For this reason, the main memory which is used to store the program which is currently being run and the data that is currently required by that program, should be stored in memory which has an access time matching processor speed sufficiently to avoid any serious slowing down of the overall processing rate. With today's technology, this means fast dynamic RAM, even faster static RAM or, faster still, memory actually on the processor chip.

The RAM that is used can, on modern processors like the 486 and Pentium, have an access time that is somewhat slower than the processor, thanks to the use of cache memory. In fact, what we have in a modern PC is a memory hierarchy with varying access speeds. At the top of the hierarchy are the handful of memory locations which comprise the processor registers, these are on the processor chip and can be accessed as fast as the processor works. Beneath that, on a chip like the 486DX2, is the primary cache memory, in this case 8KB of very fast RAM, also put on the processor chip and having access times which will not cause the processor to wait.

The next stage in the memory hierarchy is the secondary cache. On the more powerful modern systems this will probably consist of 256KB of very fast, 25ns or better access time, static RAM. We now come to the main memory of the system, the memory in which the current program and data is stored and on a modern 486 this will probably consist of about 8MB of dynamic RAM, with an access time of between 80 and 100ns.

The programs and data stored in main memory are down/uploaded from a data storage device such as a hard or floppy disk drive. The data transfer rate between main memory and disk is fast, with an IDE hard drive somewhere between 1 and 5MB per second, but nowhere near as fast as data transfer from one part of main memory to another. In order to optimise this transfer rate, another layer of memory hierarchy is often added, the hardware disk cache. This is basically a couple of megabytes of dynamic RAM

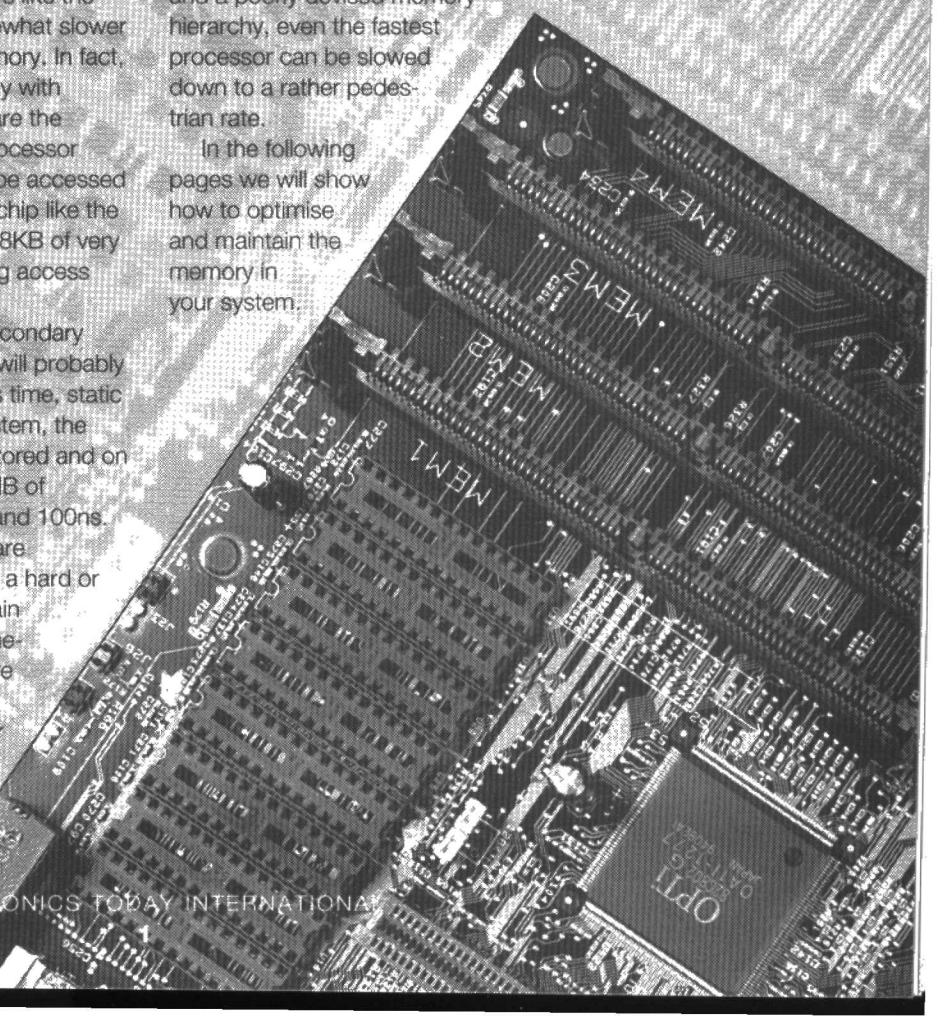
that act as a buffer between main memory and the disk and which can considerably speed up data transfer rates.

There are, of course, other much slower forms of memory, such as tape drives and CD-ROM drives. The average CD-ROM will, for example, transfer data at between 150 and 300KB per second, very slow when compared with a hard disk. By and large, these are used for archival purposes, with data being dumped from them to a hard disk and then into the main memory.

The general rule in this memory hierarchy is that fast memory is rare and expensive, whilst slow memory is common and cheap. Thus we have just a few processor registers and a small area of primary cache in which to run code at an optimum rate, a much larger area of main memory in which to run code at a slower rate, and an even larger area of hard disk, tape, CD-ROM, etc., in which to store data and programs. It is up to the programmer to make sure that these resources are used in the optimum manner.

However, it is up to the user of the system to make sure that those memory resources are available so that programs written for them can be run. If you are going to run the latest software, then you will need lots of memory and the right sort of memory. Indeed, having the right memory is more important than having the latest and fastest processor, for with the wrong memory and a poorly devised memory hierarchy, even the fastest processor can be slowed down to a rather pedestrian rate.

In the following pages we will show how to optimise and maintain the memory in your system.



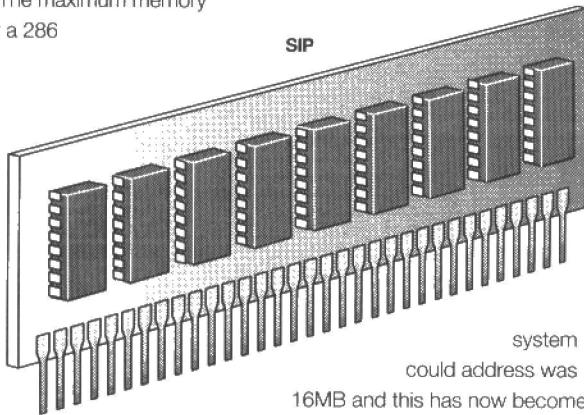
Main Memory

In all the early models of the PC, in general those produced prior to the introduction of the 386, the motherboard was designed to take either 640K or 1MB of RAM. If you wanted any more, then you had to resort to using a memory expansion board that plugged into one of the expansion sockets.

There were many reasons for this. For a start, memory chips at that time had a much smaller capacity than they do now, so fitting a large amount of memory on the motherboard would have involved using a lot of board space and been prohibitively expensive. Then there was the fact that the system and processor architecture employed in these early PCs made using more than 640K rather cumbersome.

As memory prices dropped and changes in both processor and operating system design allowed memory to be easily used above the 640K boundary, so systems started to be produced with larger amounts of RAM. A trend which was reinforced by applications software developers producing programs which needed more and more memory.

The maximum memory that a 286



system could address was

16MB and this has now become the new upper limit for memory as far as many

system and software designers are concerned. However, it should be noted that 386/486/Pentium systems can actually address a lot more memory than this, so there is no practical reason why a modern system should be limited to 16MB of RAM.

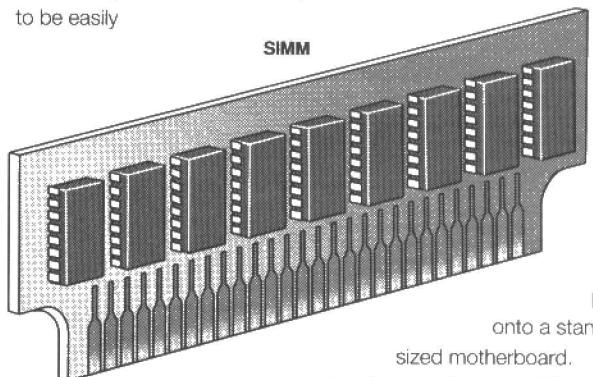
Initially, the demand for more memory was met by the use of expansion cards, however speed limitations of the expansion bus meant that as processors became faster, so the expansion memory was slowing down the whole system. This was despite the development of the much faster EISA and MCA bus systems.

The result of this limitation was the development of techniques to bypass the use of the expansion bus and directly mount large amounts of memory directly on the motherboard. Some of these early motherboards had large numbers of standard DIL chip sockets waiting to be filled. Other manufacturers introduced special proprietary sockets into

which special memory adapter boards could be fitted.

The problem with these early memory adapter boards was that they were not standard and every major manufacturer had their own design. This can mean that getting replacement memory for such systems can be virtually impossible.

Fortunately a degree of standardisation was imposed on the PC designers by the chip manufacturers. What the chip manufacturers did was to put nine 1Mbit \times 1 surface mount memory chips onto a small module, to create what was in effect a single, much larger 1M \times 9 memory chip (the reason why nine bits are used is discussed on page 38). These small memory cards are mounted vertically in special sockets and because they are compact, they allow a large amount of memory to be easily

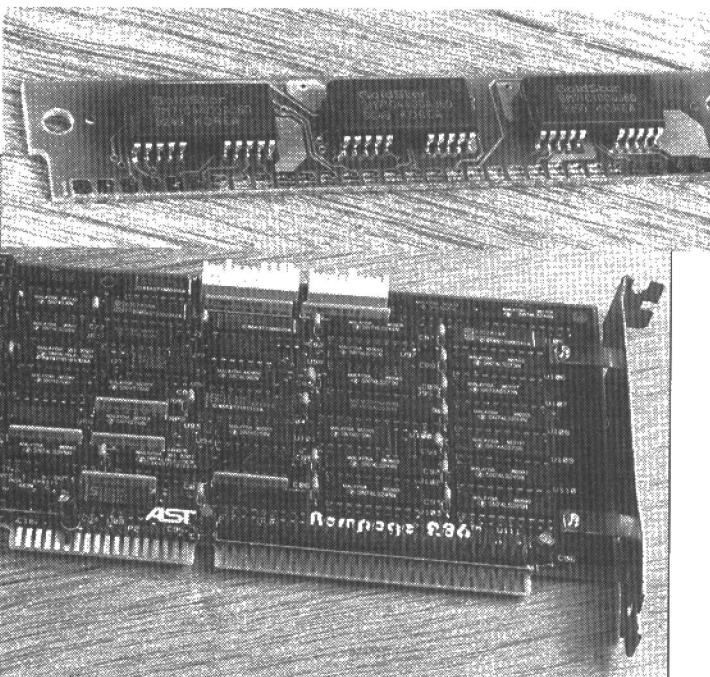


placed onto a standard sized motherboard.

Now there are just two different types of memory modules, the SIP (Single Inline Package) and the SIMM (Single In-line Memory Module). They are now available in either 1M \times 9 or 4M \times 9 format, making it possible on some motherboards to have up to 64MB of RAM actually on the motherboard.

Because SIMMs use sockets rather than the pin connections of SIPs, there is a tendency for manufacturers to go for SIMM based systems since they are easier to fit and there is less danger of bending or breaking a pin. When adding more memory or replacing memory it is important to ensure that the right type of module is purchased. Although if you already have some memory modules but they are the wrong sort, don't despair. Special converter modules are available through some specialist suppliers.

Of course, many systems do not exclusively use either standard DIL packaged memory chips or SIP/SIMM modules, some use a combination of both. Neither is the standard SIP/SIMM necessarily nine bits wide, some are only eight bits wide. There is a lot of variation.



Memory Organisation

There is an enormous variation in the types and specifications of memory chips and modules that are used in PCs. Indeed, any given chip or module can vary in three main parameters - for a start there is the size of the chip, the number of bits of memory that it contains, 64K, 256K, 1M, or even 4M. Then there is its organisation, thus a 1Mbit chip could be organised as 1M x 1, or as 256K x 4, or even as 128K x 8, where the 'x 8' indicates the fact that each memory address on the chip contains eight bits of memory. Finally, there is the speed of the chip, this is the time taken to access memory stored on the chip, it is important that the access time matches as closely as possible the speed of the processor. If it is significantly slower, then the processor will have to wait and processing speed will be seriously reduced.

In normal use, it is unnecessary to know what type of memory is used and how it is organised, but if you want to add more memory, or replace faulty memory, then this kind of knowledge is essential. In most cases this information should be provided in the manual, but in its absence one can usually work out the memory type and organisation without too much trouble.

The first thing to check is memory organisation. In any PC, the smallest data bus is eight bits wide and this means that the memory has to be organised to provide eight bits of data for every memory address. To do this the system designers could use eight memory chips which are each capable of providing one bit of memory for each address and organised as shown in Figure 2. Conversely, they could use two memory chips, where each chip provides four bits for each address, as shown in Figure 3. It should be born in mind, however, that the use of a parity bit - see page 38 - means that in reality nine bits are used rather than eight. The two drawings on this page are shown with only eight bits for the sake of clarity.

With SIMMs and SIPs, there is no need to worry about whether one bit or four bit, or even eight bit, wide memory chips are used, since the modules are always either eight or nine bits wide, but with the standard DRAM chips found on older machines, it is important to use the right chips when upgrading memory or replacing faulty memory. The easiest way of doing this is to check the type of chips which are currently installed. To help with this task, a table of the chip markings for some of the commonest types of memory chip produced by the major manufacturers is shown on this page.

Also included is a table of memory speeds - Table 2 - with relation to the speed of the processor. When installing extra memory it is essential that the memory you use has a quoted access time which is sufficiently fast for the processor. Thus, for example, if you are upgrading the memory on a system with a 20MHz system clock, then it is essential that the motherboard memory that is used should have an access time of 80ns or better.

Expansion memory added via one of the expansion slots can be somewhat slower, because the expansion bus itself has a maximum access speed which will, as we have seen, unavoidably slow down the processing speed of any system running at more than about 8MHz.

This table of memory access speed requirements is only approximate. Indeed, some machines seem to work with memory which in theory should be too slow, but it should be noted that such systems can as a result of speed problems generate odd errors when they warm up. Finally, if you have any doubts about what speed memory to use, then use the next highest speed, adding memory that is too fast for the system will do no harm, it will simply cost a little bit more.

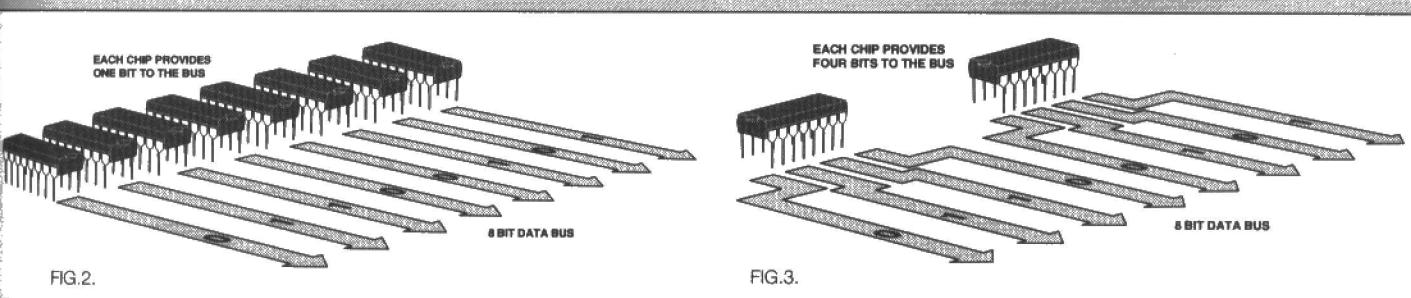


Table 1.
Minimum memory speed with relation to system clock:

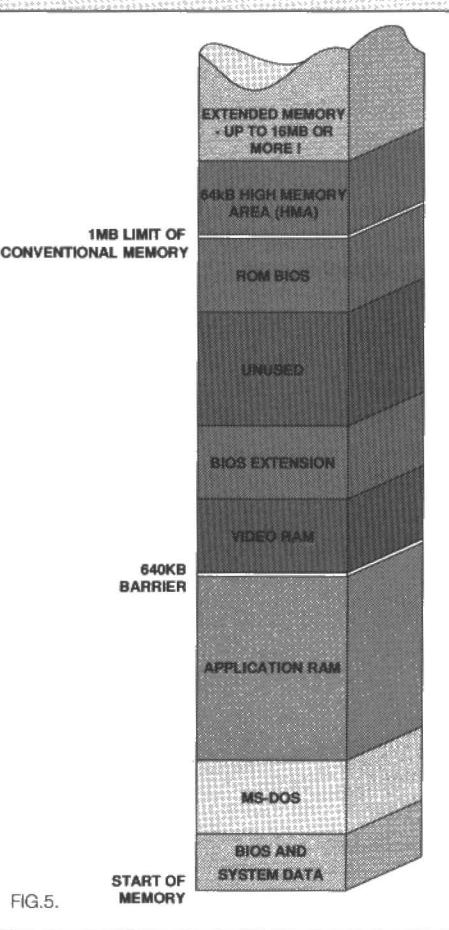
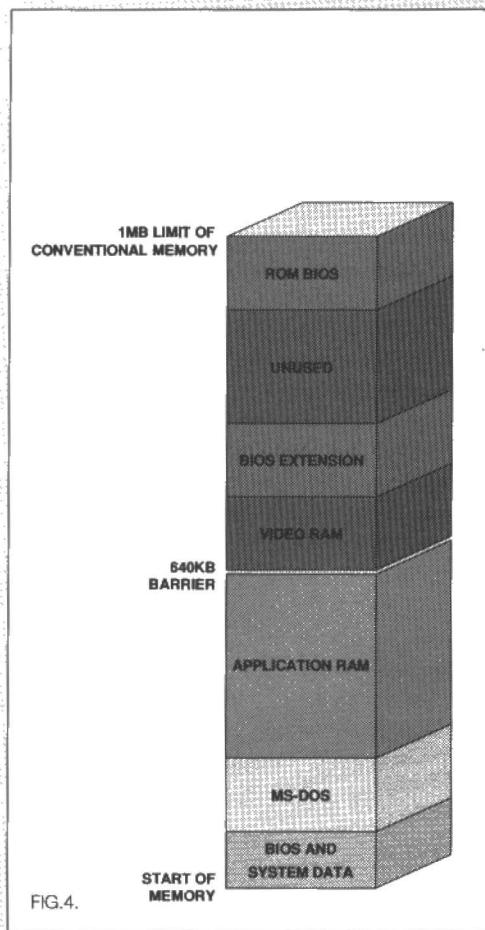
| System Clock | Motherboard RAM speed in nanoseconds | Expansion RAM speed in nanoseconds | Chip speed designation on motherboard RAM |
|--------------|--------------------------------------|------------------------------------|---|
| 8MHz | 150ns | 150ns | 15 |
| 10MHz | 120ns | 150ns | 12 |
| 12MHz | 100ns | 120ns | 10 |
| 16MHz | 100ns | 120ns | 10 |
| 20MHz | 80ns | 100ns | 80 |
| 25MHz | 80ns | 100ns | 80 |
| 33MHz | 70ns | 100ns | 70 |

Table 2.
Types of dynamic RAM chip which might be found in a PC:

| | 64Kx1 | 64Kx4 | 256Kx1 | 256Kx4 | 1Mx1 |
|------------------------|---------|----------|----------|-----------|---------------|
| Generic | 4164 | 41464 | 41256 | 414256 | 411000/511000 |
| Fujitsu | MB8264A | MB81464 | MB81256 | | MB81C1000 |
| Hitachi | HM4864 | HM50464 | HM51256 | HM514256 | HM511000 |
| Hyundai | HY5164 | HY53C464 | HY51C256 | | |
| IBM | 4164 | 41464 | 41256 | 414256 | 511000 |
| Intel | 2164 | | | | |
| Micron Tech | MT4264 | | MT1256 | MT4C4256 | MT4C1024 |
| Mitsubishi | M5K4164 | | M5M4256 | | M5M4C1000 |
| Mostek | 4564 | | | | |
| Motorola | MCM6665 | | | | |
| National Semiconductor | NMC4164 | | | | |
| NEC | uPD4164 | uPD41464 | uPD41256 | uPD414256 | uPD421000 |
| OKI | M3764 | M41464 | M41256 | M514256 | |
| Panasonic | MN4164 | | MN41256 | | |
| Samsung | KM4164 | KM41464 | KM41256 | KM414256 | KM41C1000 |
| Siemens | HY4164 | HY41464 | HY41256 | | |
| Texas Instruments | TMS4164 | TMS4464 | TMS4256 | TMS44256 | TMS4C1024 |
| Toshiba | TMM4164 | TMM41464 | TMM41256 | | TC511000 |

The PC Memory Map

When upgrading a system, it is helpful to know how the PC actually uses memory and what programs and data are stored in specific areas. Unfortunately, for historical reasons, the memory map for the PC is rather complicated and as a result, we have to divide memory into three specific types.



There is conventional memory, extended memory and expanded memory.

Conventional memory is limited to a maximum of 1MB and extends from address 0 through to 1024K (0000 through 9FFF hex). Of this 1MB, the bottom 640K is conventionally set aside for programs and data, whilst the 384K block of memory between addresses 640K and 1023K (A000 through FFFF hex) is reserved for use by peripherals and device drivers, and is known as 'Upper Memory'.

From a system point of view it is this upper memory area (or UMA) which is most interesting.

Looking at Figure 4, one can see that it is broken down into four main areas. At the bottom is an area extending from A000 to BFFF hex, which is allocated to the video adapter and is used to store a bit image of the high resolution display.

Above the video is an area of memory extending upwards from C000. This is used to store the program code and data for the system's installed device drivers and BIOS extensions. Some of these drivers have standard memory areas allocated to them, others are simply put into the next available section of free memory. In fact, the device drivers on most systems will only occupy memory up to about CFFF, thus leaving the memory area between D000 and EFFF hex unused.

Another very important application for the UMA is to map ROM into RAM and in a great many PCs, the ROM BIOS is mapped into the memory area F000-FFFF. This does not mean that the contents of the BIOS ROM is copied into RAM in this memory area, what it usually means is simply that the BIOS ROM chip has been assigned this area of memory.

Of course, the BIOS ROM may not be the only ROM in a system. There are, in all probability, some special BIOS extension ROMs, such as the Video BIOS and a Hard Disk BIOS. Where these exist, they will also have to be assigned blocks of memory in the UMA. Network cards are another case where special memory

Fault Finding in Memory

When any PC powers up, it goes through a special set of BIOS routines known as the Power On Self Test, or POST, routines. An important element of these routines is checking memory and in fact BIOS displays the amount of memory checked as a rolling upward figure on the screen as it is doing it. These POST tests will find nearly all permanent chip errors and will identify the location of the memory bank where the error was located.

The POST routines are the main way of checking memory, although there are, in addition, a number of memory exercising routines which will perform more complex error checking and can prove useful when attempting to locate a heat sensitive chip. But there is another type of memory error which is not so easy for the POST routines to detect, the transient error. It is a type of error which IBM PC/AT systems and 100% compatibles are able to detect thanks to bit nine of memory.

Anyone who is in the least bit familiar with computers will know that a Byte consists of eight bits and it may therefore seem rather strange that the PC uses memory that is nine bits wide as opposed to eight. The answer lies with IBM, the original designers of the PC. On large systems, it had been standard procedure to use memory that was nine bits wide and employ the extra bit to perform an error checking function on the other eight bits. It was what is known as the parity bit. Because IBM's designers were used to systems that had memory with a parity bit, they built it into the PC and it has stayed ever since.

The value of having this ninth parity bit is that it enables the system to automatically detect transient errors. These are temporary errors that can occur in memory as a result of something like a spike in the mains, or more exotically, a high energy alpha particle derived from a cosmic ray hitting your PC.

However, memory chips today are very reliable and cosmic rays and power spikes fortunately, rather rare, so there is a good argument for saying that there is no need for the parity bit. It does, after all, add a lot of extra unused memory and therefore cost, to a system. Indeed, a handful of systems are now being produced with eight bit memory and no parity bit. This could become much more common in the future and it is worth checking before you buy any extra memory whether your system has eight bit or nine bit memory.

blocks will have to be assigned in the UMA

The memory area above the 1024K boundary is referred to as extended memory, so if your PC has 8MB of RAM, then we can divide it into 7MB of extended memory and 1MB of conventional memory. As is shown in Figure 5, extended memory can go up to 16MB or even more.

You will only find extended memory in 286 and later systems, the earlier 8086 and 8088 based systems were unable to make use of memory above the 1MB boundary. This

of view, it is very easy to use.

Infinitely more difficult to use is the expanded memory systems employed on many older PCs. This system was developed to overcome the 1MB addressing limitations of DOS, but unlike the memory manager used in extended memory, the expanded memory manager does not allow DOS to directly access more than 1MB. Instead, as shown in Figure 6, it sets up a 64KB page frame in the UMA and transfers the contents of expanded memory to that page frame as required.

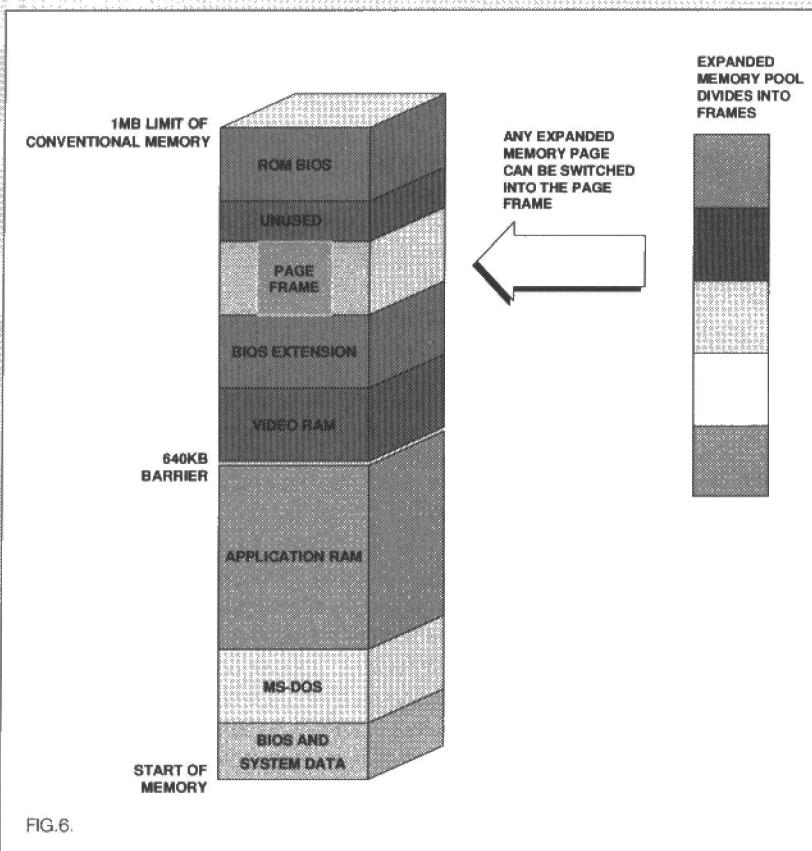


FIG. 6.

was because they could not run in protected mode, only in real mode and also because they lacked facilities for handling memory above 1MB.

Indeed, DOS itself cannot address memory above the 1MB boundary, and a memory manager is therefore used to provide access to extended memory. In DOS, the standard extended memory manager is the HIMEM.SYS driver. In operation, a DOS program written to access extended memory does so by requesting access from the memory manager.

It is the memory manager which allocates memory to an application and then returns it to the pool of available memory when the program has finished with it. The memory that it allocates is always in the form of a contiguous block, which means that from a programming point

of view, it is very easy to use. Expanded memory is primarily used when memory resides on a memory expansion card. Essentially, what the expanded memory manager is doing is mapping a 64KB chunk of the expansion card memory, which could have up to 2MB of RAM on it, into the UMA where DOS can use it as it would use any part of conventional memory. However, the fact that direct access is limited to 64K at a time, makes expanded memory rather cumbersome to use and will give much slower performance.

In general, we can conclude by saying that if anyone wants to expand the memory in a system, it is far better to use extended memory and, if possible, memory that can be added to the main motherboard. Expanded memory should only be used as a last resort.

It should be noted that in this issue, we are only looking at main memory in a PC, there are in addition several other sorts of memory and applications for memory. Last month, we briefly looked at cache memory and the BIOS ROM. In a couple of months' time, we will be looking at the video memory that is on the adapter card of many VGA and SVGA display cards. This will be followed by a look at disk caching memory and printer buffer memory.

Anti-static Precautions

We all tend to build up a small static charge in our bodies as a result of moving around, the friction of clothes, etc. Most of us are familiar with that childhood trick of rubbing a comb through our hair and then using it to pick up a small scrap of paper. Normally, such charges are small and quickly dissipated. However, the widespread use of synthetic fibres in clothing, carpets, and furnishing fabrics, plus the use of synthetic rubber soled shoes, means that there are fewer chances for the static charge to dissipate. Consequently the human body, acting as a large capacitor, can store quite a considerable charge.

Static electricity is an ever present hazard when handling electronic components, because the very high voltages present when a build up of static electricity in the body of the handler is discharged through a device containing sensitive integrated circuits, can easily destroy those circuits.

The best way to eliminate this problem is to discharge the static from your body prior to handling any sensitive circuitry. This can be done by simply touching something conductive which is attached to earth, for example a mains water pipe, a metal central heating radiator, or electrical earth on the mains system.

The trouble is that it is not always convenient to have to keep touching an earthed metal object in order to discharge your static electricity. One better way is to use an anti-static wrist strap. This is a simple conductive band held onto the wrist by a Velcro fastener. The conductive strip is attached to a cable via a 1Mohm resistor for safety earth. This simple little wrist strap will drain away all damaging static charge, although be careful to remove the strap before working with live voltages.

Anti-static wrist straps are available from Maplin Electronics and cost £6.95 each.

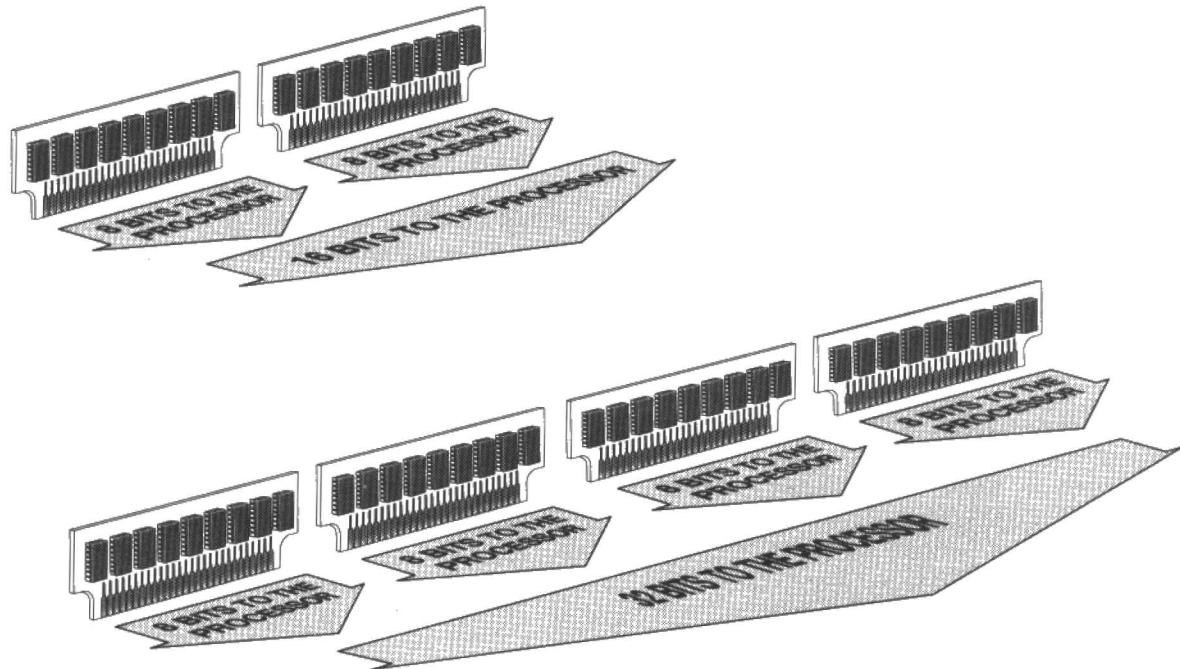


FIG.7.

The Use of Memory Banks

The memory on a PC is, as we have already seen, organised as eight bit, or byte wide format plus in most cases an extra ninth, or parity bit. As a result the majority of SIMM or SIP modules found in PCs today are configured as $1\text{M} \times 9$, with an increasing occurrence of the much larger $4\text{M} \times 9$ modules.

This means, of course, that memory tends to be organised in multiples of such blocks, so in a modern system, memory will increase in multiples of 1M , or even 4M . In older machines the increments were much smaller. Thus, an XT system would probably have had its memory organised as $256\text{K} \times 9 + 256\text{K} \times 9 + 64\text{K} \times 9 + 64\text{K} \times 9$, to give a total memory size of 640KB . Here the increment was just 64K bytes at a time.

The smallest increment in memory size is often referred to as a 'bank', thus in the above example the XT has its memory organised in banks of 64K and in the modern system the memory is organised in banks of 1MB and 4MB respectively.

So far we have assumed that a bank is a block of memory just eight or nine bits wide and this is certainly true on an XT system which has an eight bit wide data bus, so to add a block of memory to such a system you will need to use nine one bit chips. However, later systems have wider data buses and therefore wider memory blocks. Thus 286 and 386SX systems have a 16-bit data bus and a block of memory in such systems will consist of eighteen one bit chips. Full 386 systems and all 486 systems, have an even wider 32-bit data bus and therefore require 36 one bit chips in a block. This is shown diagrammatically in Figure 7, although it should be noted that to make the diagram clearer the parity bits have not been included.

The arrangement of memory in banks is also the reason why, in some situations, it is necessary to remove memory in order to further upgrade the amount of memory in a system. To the layman this always sounds totally illogical, why remove some chips or modules and then replace them with others?

The reason is that many systems allow memory chips or modules of different sizes to be used in the same sockets. It is, for example, common for a machine to be able to take either $256\text{K} \times 9$ or $1\text{M} \times 9$ SIMMs in the same socket. If such a system allows a maximum of eight memory modules, then it could have a maximum memory size of $256 \times 8\text{K}$ or 2MB when fitted with 256K modules, while increasing the module size to 1M will increase the maximum memory size to 8MB . This is

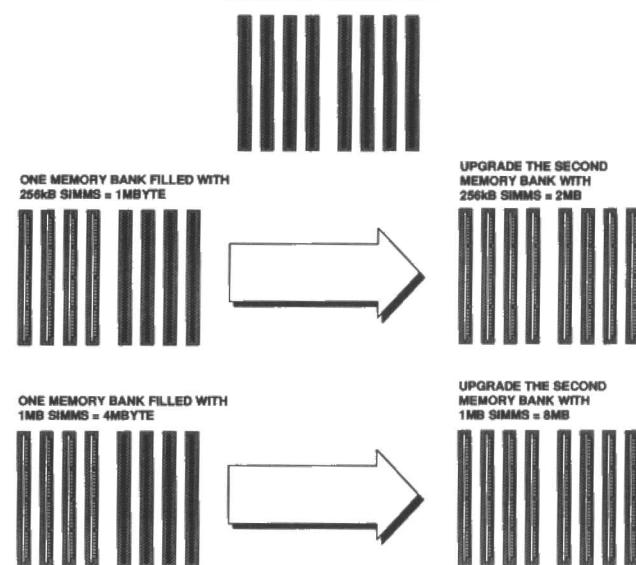
shown diagrammatically in Figure 8. The module size is often defined by a jumper or switch on the motherboard.

It should be noted that the diagram in Figure 8 is for a 486 system with a 32-bit data bus. This means that, although there are eight available SIMM sockets, there are actually only two memory banks, each consisting of four SIMMs. So in reality although there may be 8MB of memory in the system, it is actually organised as $2\text{M} \times 36$. It is important to remember this since, if you are upgrading memory, it is no good simply buying one SIMM and filling one socket, if you have a 32-bit system then you have to buy a minimum of four SIMMs. In other words, enough memory to fill a complete bank.

It is always a good idea when buying a new machine to opt for one that has the maximum upgrade potential, being able to add more memory is often more important than adding a faster processor. To achieve the maximum upgrade potential means having a system which uses the largest size SIMMs that are available, so that it uses as few complete memory banks as possible.

TWO MEMORY BANKS FREE = 0KB

FIG.8.



Upgrading Memory

Before attempting to upgrade the memory in your system and certainly before buying any memory, it is important to look at the organisation of the existing memory in your system, particularly the question of memory banks. After careful analysis, you may well find that it is impossible to add further memory to the motherboard, in which case there are two options. Either use an extended memory card that plugs into one of the expansion bus ports, or buy a new motherboard which can take more memory.

In the long run, you will probably find that purchasing a new motherboard is cheaper and results in fewer problems (one can avoid having to use slow extended memory, for a start). With a careful choice of motherboard, you should be able to recycle your existing processor and perhaps even your memory chips, thereby cutting down the cost.

The second possible result of your analysis is that there is no room to add further memory, because all the banks are full of small sized chips or modules (assuming that your system will take either large or small sized chips/modules in the same sockets). In this kind of situation, if you want more memory then the only option is to remove the memory currently in your system and replace it with larger size chips or modules.

If you remove the chips or modules carefully, it is often possible to recoup some of the cost of the new memory by selling them to other users who are looking for a cheap way to expand their system. Don't forget to switch off and disconnect the system before opening the case and discharge any static before touching the chips or modules. When removing chips, it is a good idea to use a proper chip extraction tool.

Of course, your analysis of the system may well reveal that it can be expanded by simply adding further chips or modules to form another memory bank. Where chips or modules have to be bought, either for replacement of, repair of, or in addition to existing memory, then it can not be over emphasised how important it is that the specifications are correct, particularly with regard to access speed.

Whenever you add memory chips to a system, first switch off and disconnect from the mains before opening the case, then carefully discharge your body's static before touching any of the chips or the motherboard. Keep discharging your static at frequent intervals.

If you are adding DIL chips, then the most frequent source of problems is a bent or in some cases a broken pin. You will stand a better chance of success if you carefully straighten the pins before attempting to insert them into the socket, as shown in Figure 9. You can do this with a small pair of pointed nose pliers or better still with a proper chip insertion tool. These are quite cheap and a good investment if you intend to insert a lot of chips.

However, before attempting to insert the chip it is very important that it is inserted the right way round. Chips usually have a notch at one end, or a small depression in one corner and this should be matched up with the corresponding markings on the PCB, as shown in Figure 10.

On boards which allow two different sizes of chip to be used, there is often the added complication of a two sized chip socket. You will commonly find 18/16 pin sockets to take either a 16 pin chip or an 18 pin chip, and 18/20 pin sockets for 18 and 20 pin chips. To use these sockets simply select the type of chip you want to use and align pin 1 correctly.

The development of SIMM and SIP modules has made it much easier to add or remove memory, particularly in the case of SIMM chips. A SIP socket is just a long line of holes into which the pins on the SIP module fit, installing a module is simply a matter of lining up the pins with the appropriate holes and carefully pushing them in, making sure that no pins are bent. But before doing this check two things - firstly that the module is the right size. There are 30 pin SIPs and 32 pin SIPs and you must have the right type. Secondly you must make sure that the SIP is inserted the right way round and you will find that a notch on the SIP marks pin 1. This is shown in Figure 11.

A SIMM module has no pins, instead it looks a little like a miniature PCB with edge connector pads along the bottom. This type of module can be inserted by simply pushing it into the connector at an angle and then carefully moving the SIMM to an upright position so that it locks into place, see Figure 12. Once again, it is important to make sure that the SIMM is inserted the right way round, as shown in Figure 13.

The above dissertation on upgrading the memory in a system also largely applies to replacing faulty memory in a system. The POST routines should help in locating the memory bank with the fault, and some careful swapping of modules should enable one to pinpoint the actual faulty module. Fortunately, memory faults are very rare today and SIMM/SIP modules are much more reliable than earlier DIL memory chips. However, because SIMM/SIP modules are constructed using surface mount technology it is not really worthwhile trying to repair modules. Just consider a module as being the same as a chip, throw it away and replace it with a new one.

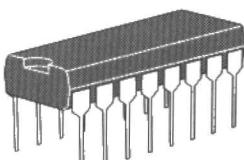
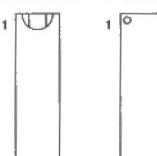


FIG.9A. AFTER STRAIGHTENING

THE 'PIN ONE' END OF A CHIP IS MARKED BY A SEMI-CIRCLE OR A DOT



THE 'PIN ONE' END OF A SOCKET IS MARKED BY A SEMI-CIRCLE, NOTCH, OR DOT ON THE MOTHERBOARD

FIG.10.

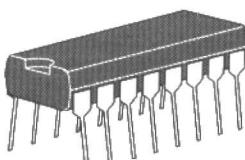


FIG.9B. BEFORE STRAIGHTENING

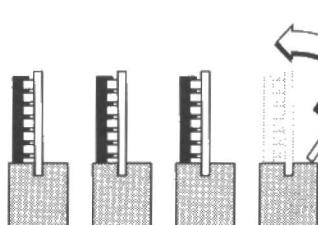


FIG.12.

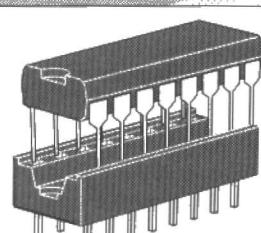


FIG.9C.

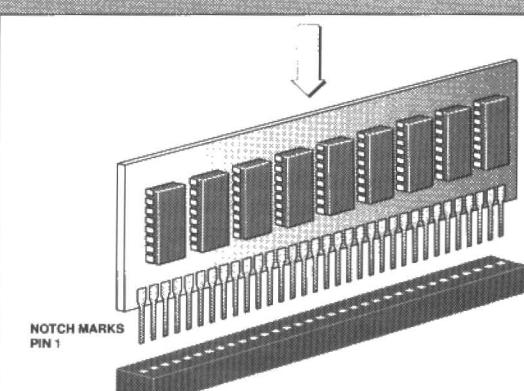


FIG.11.

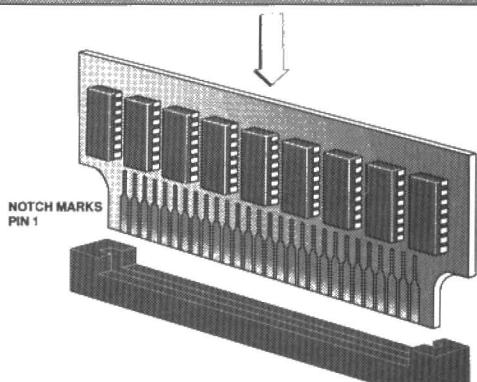


FIG.13.

PCMCIA

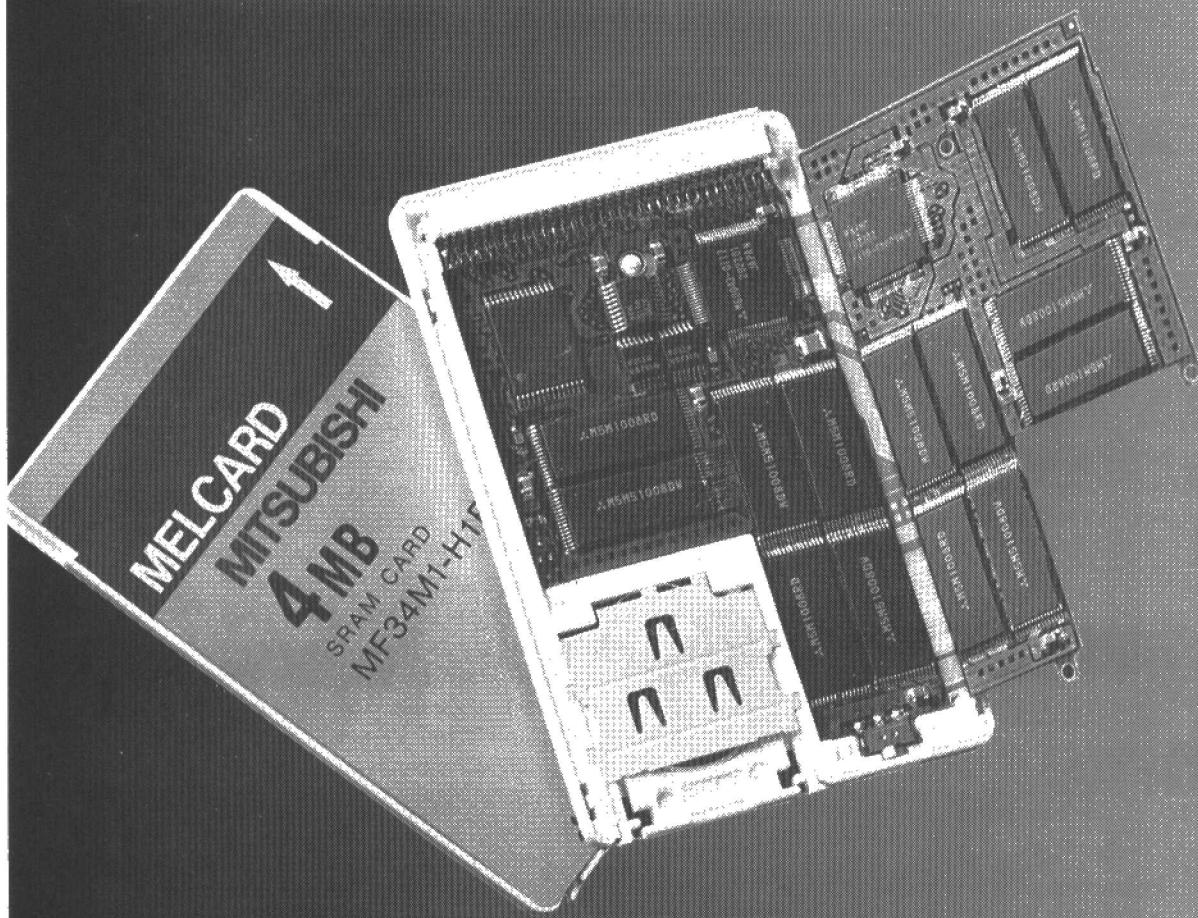
Small portable PCs, in particular the very small notebook machines, have one major drawback. They have no room for fitting any expansion hardware. They cannot take even a single adapter card and in many cases cannot be given more memory. To overcome this problem, manufacturers have developed the PCMCIA connector. This makes use of credit card sized expansion cards that can simply slot into the connector at the side or rear of the machine.

The first PCMCIA cards to appear were plug in modems, an obvious add-on for notebook and portable PCs. These were closely followed by memory cards, which are now available with capacities of up to 4MB. More recently, PCMCIA cards with a wide range of different functions have started to come onto the market. There are data acquisition cards with a full range of analogue and digital I/O that can be used to turn a notebook PC into a portable data acquisition system. There are also IEEE 488 bus interface cards to enable a notebook to control instrumentation. There is even a PCMCIA global positioning system card, which allows a portable computer to be used as a navigation system.

Most of the PCMCIA memory cards have memory back-up built into the card which allows them to retain data stored in them for a reasonable length of time. This, coupled with their fairly high data capacity, has led to such cards being used as a means of transferring data and programs between machines. The result is that an increasing number of desktop and tower machines are now being fitted with PCMCIA connectors.

I expect it is only a matter of time before we see such connectors on a wide range of intelligent equipment. The PCMCIA card is after all a very closely allied product to the smartcard. Some readers might like to experiment with these cards and their connectors and they could be a good way of transferring data between a microprocessor system, such as ETI's FORTH, Z-80 and Transputer cards, and a PC. For these readers' benefit we are including the standard pin connections for a PCMCIA card. PCMCIA cards are now widely available from many sources, although the connectors are less widely available, but can be obtained from Radio Spares.

| | |
|------|----|
| GND | 1 |
| D3 | 2 |
| D4 | 3 |
| D5 | 4 |
| D6 | 5 |
| D7 | 6 |
| /CE1 | 7 |
| A10 | 8 |
| /OE | 9 |
| A11 | 10 |
| A9 | 11 |
| A8 | 12 |
| A13 | 13 |
| A14 | 14 |
| /WE | 15 |
| NC | 16 |
| Vcc | 17 |
| NC | 18 |
| A16 | 19 |
| A15 | 20 |
| A12 | 21 |
| A7 | 22 |
| A6 | 23 |
| A5 | 24 |
| A4 | 25 |
| A3 | 26 |
| A2 | 27 |
| A1 | 28 |
| A0 | 29 |
| D0 | 30 |
| D1 | 31 |
| D2 | 32 |
| WP | 33 |
| GND | 34 |
| GND | 35 |
| /CD1 | 36 |
| D11 | 37 |
| D12 | 38 |
| D13 | 39 |
| D14 | 40 |
| /CE2 | 41 |
| NC | 42 |
| NC | 43 |
| NC | 44 |
| NC | 45 |
| A17 | 46 |
| A18 | 47 |
| A19 | 48 |
| A20 | 49 |
| NC | 50 |
| Vcc | 51 |
| NC | 52 |
| NC | 53 |
| NC | 54 |
| NC | 55 |
| NC | 56 |
| NC | 57 |
| NC | 58 |
| NC | 59 |
| NC | 60 |
| /REG | 61 |
| BVD2 | 62 |
| BVD1 | 63 |
| D8 | 64 |
| D9 | 65 |
| D10 | 66 |
| /CD2 | 67 |
| GND | 68 |



Interconnecting to a SIMM

Although it is not really worthwhile trying to repair a SIMM module, it might be worth looking at using them for other applications. They are a relatively cheap form of memory (particularly if you can get some second-hand). So, for the experimenters, here is a listing of the pin-outs on a 1M x 9 SIMM. The only difference between this and an 8-bit SIMM is the fact that the lines associated with bit nine are not connected.

Next month...

We will be looking at I/O circuitry, keyboards and pointers (i.e. mice).

| | |
|------|----|
| Vcc | 1 |
| CAS | 2 |
| DQ1 | 3 |
| A0 | 4 |
| A1 | 5 |
| DQ2 | 6 |
| A2 | 7 |
| A3 | 8 |
| Vss | 9 |
| DQ3 | 10 |
| A4 | 11 |
| A5 | 12 |
| DQ4 | 13 |
| A6 | 14 |
| A7 | 15 |
| DQ5 | 16 |
| A8 | 17 |
| A9 | 18 |
| NC | 19 |
| DQ6 | 20 |
| W | 21 |
| Vss | 22 |
| DQ7 | 23 |
| PRD | 24 |
| DQ8 | 25 |
| Q9 | 26 |
| RAS | 27 |
| CAS9 | 28 |
| D9 | 29 |
| Vcc | 30 |

TOW GUARD

Prevent thieves towing away your caravan with Terry Balbirnie's anti-theft device

Caravan crime, like all vehicle crime, has become very widespread. Although petty theft of caravan contents has always been a problem, there appears to be a trend towards the theft of caravans themselves. By removing your trailer, the thief can rifle it of any expensive equipment at his leisure then sell the caravan. Whereas many people now protect their home with some type of intruder alarm, the caravan is often left very vulnerable and makes an easy target for the thief.

There are several types of caravan alarm and security device available. A wheel clamp is particularly effective, although it needs to be of a good quality and this makes it expensive. The interior of the caravan can be protected using a PIR sensor or an alarm operated by door and window magnetic reed switches.

Tow Guard is not meant to be the sole source of protection. It is designed to complement any other measures the owner decides to take. Even so, it should be remembered that no anti-theft system is entirely foolproof. The best that can be achieved is to provide time so that the thief can be apprehended or will run

away before stealing any goods or causing damage.

Tow Guard operates a loud siren - the familiar yelping type associated with car alarms - the moment the caravan plug is slightly inserted into the tow-car's electrical socket. It will then operate for some preset time between 15 seconds and 3 minutes even when the plug is removed.

No thief will risk proceeding with an alarm sounding and it would take too long to disable it. Also, he has no idea whether it will eventually stop. The chances are, he will beat a hasty retreat. If he had been aware of the type of alarm - which is unlikely - he could have risked driving without coupling the electrics. On the other hand, he will not want to attract suspicion by having no brake lights or flashing indicators. The last thing he wants is to be stopped by the police.

Windproof

This system has several advantages. Firstly, it is inexpensive to construct yet provides good security. Also, there is no need for entry and exit time delays which are used in some systems. The owner can therefore enter his own caravan without taking any special precautions. A

further point is that the alarm will not trigger in strong winds as sometimes happens with units using motion sensors. The alarm could also be arranged to operate using microswitches and these could be fitted to corner steadies, gas bottle covers, etc. if required. More will be said about this later.

The unit is built in a small plastic box and this will be situated in a convenient locker. The siren may be mounted under the caravan. Alternatively, it could be sited in the same locker as the main unit and a few holes made in the floor for the sound to pass through.

The standby current drain is only around 100mA and 300mA with the alarm actually sounding. The integral battery pack should easily provide one year's service, even if the alarm has sounded a few times.

Note that there are two sockets on a modern outfit and the one referred to is the 12N (normal) socket - that is, the one which feeds the caravan road lights. Tow Guard is triggered when pins 2 and 3 touch the corresponding ones in the drawbar socket. Pin 2 is normally responsible for operating the rear fog lights and pin 3 connects the road light

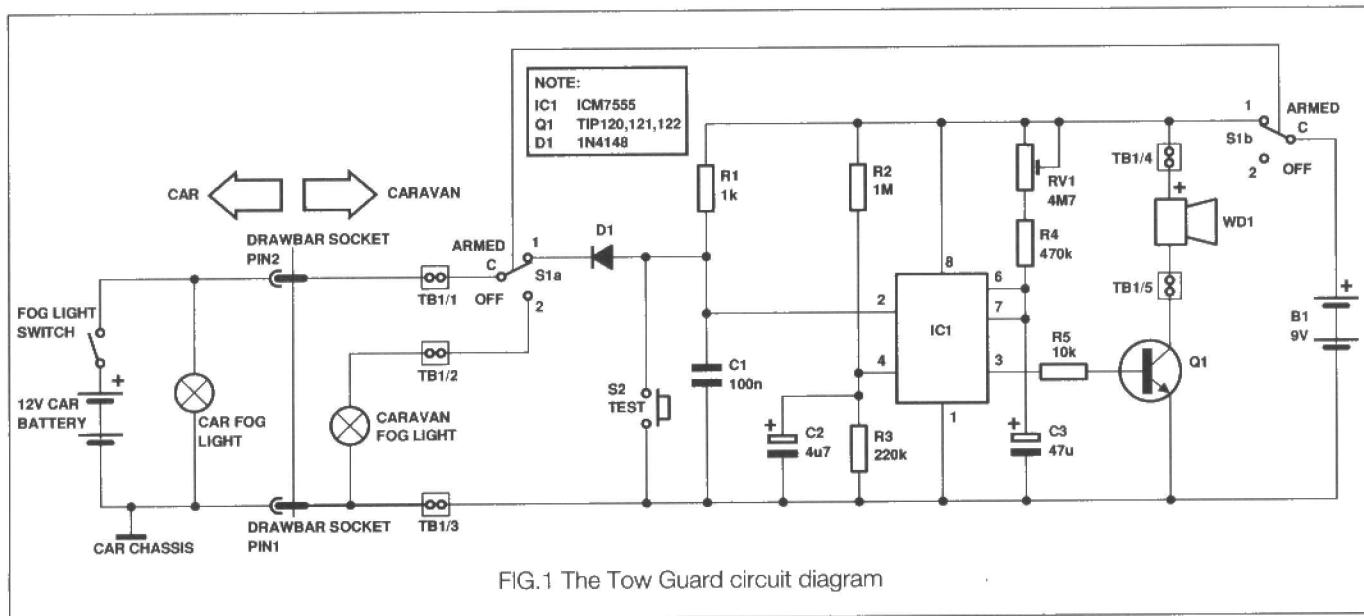


FIG.1 The Tow Guard circuit diagram

common wire (earth) to the car chassis. While the alarm is armed, the feed wire from pin 2 is switched over to its new purpose. It is impossible to drive away without switching it back again because the user would, of course, set off the alarm himself.

Circuit description

Figure 1 shows the complete circuit, with the appropriate pin connections between towing vehicle and caravan when the 12N plug is inserted. The part to the left of the dotted line shows the simplified wiring to the existing car fog lamps.

In the main unit, the wire leading from pin 2 (fog light feed) is connected to the common contact (c) of double-pole switch, S1 pole a. While off, this connects to the caravan fog lights so these will operate normally. Also, pole b disconnects the internal battery, B1, so no current is drawn. With S1 in the armed position, pole b establishes the supply and pole a connects the fog light feed wire to diode, D1, cathode instead of the lights.

With the alarm armed and with the caravan idle, the connection to D1 leads nowhere (because the caravan plug is disconnected) and nothing happens. When the caravan plug is inserted, pin 3 connects the car chassis (earth) to the road light common wire. At the same time, pin 2 transfers a low state (battery negative voltage) via the fog light filaments on the towing car to D1. This low state is applied to IC1 trigger input, pin 2.

Monostable

Integrated circuit timer IC1, is connected as a monostable. When triggered by the low state referred to above, the output (pin 3), goes high (positive supply voltage) for a certain time then reverts to low. The time interval depends on the values of fixed resistor, R4, preset potentiometer VR1 and capacitor, C3. With the values specified, the maximum timing is 3 minutes and the minimum, 15 seconds. Note that these times are very approximate because C3 is subject to a wide tolerance. It often turns out that the timings are considerably more than theory predicts.

While IC1 output is high, current enters the base of Darlington transistor, Q1, through current-limiting resistor R5 and the audible warning device, WD1, in the collector circuit sounds.

The reset input, pin 4, of IC1 is kept normally high via resistor R2 and this enables the device. However, when the supply is first connected, capacitor C2

charges through R2 and the voltage across it rises from zero to about 1.6V. This keeps IC1 pin 4 low for a short while and prevents self-triggering. When switched off, R3 allows C2 to discharge rapidly ready for further operation. Switch S2 is used for testing the alarm. When operated, it makes IC1 pin 2 go low and trigger the monostable. While in standby mode, the trigger input is kept high via resistor R1 and this prevents false operation.

The alarm on/off switch, S1, was a key-operated type in the prototype unit. This provides the best security. However, an ordinary switch could be used instead because it is unlikely that the thief will have time to break into the caravan to switch it off. This would save costs but whatever the type, the switch contacts

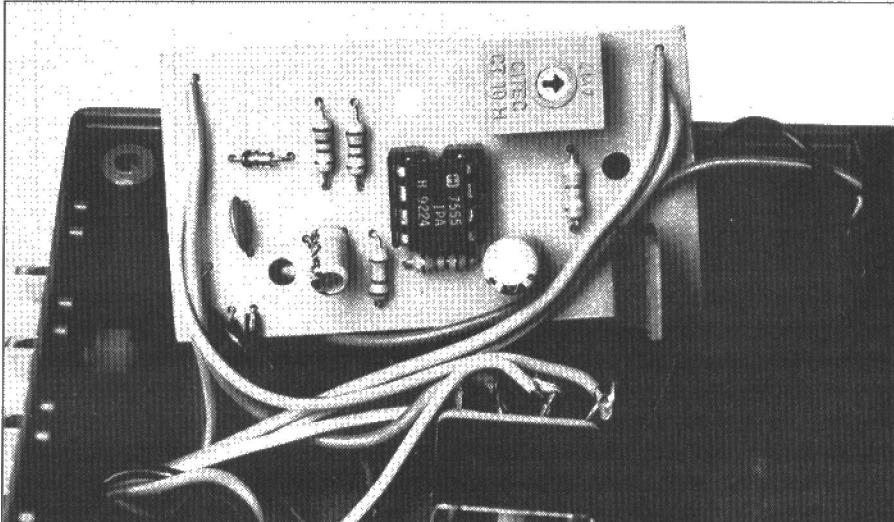
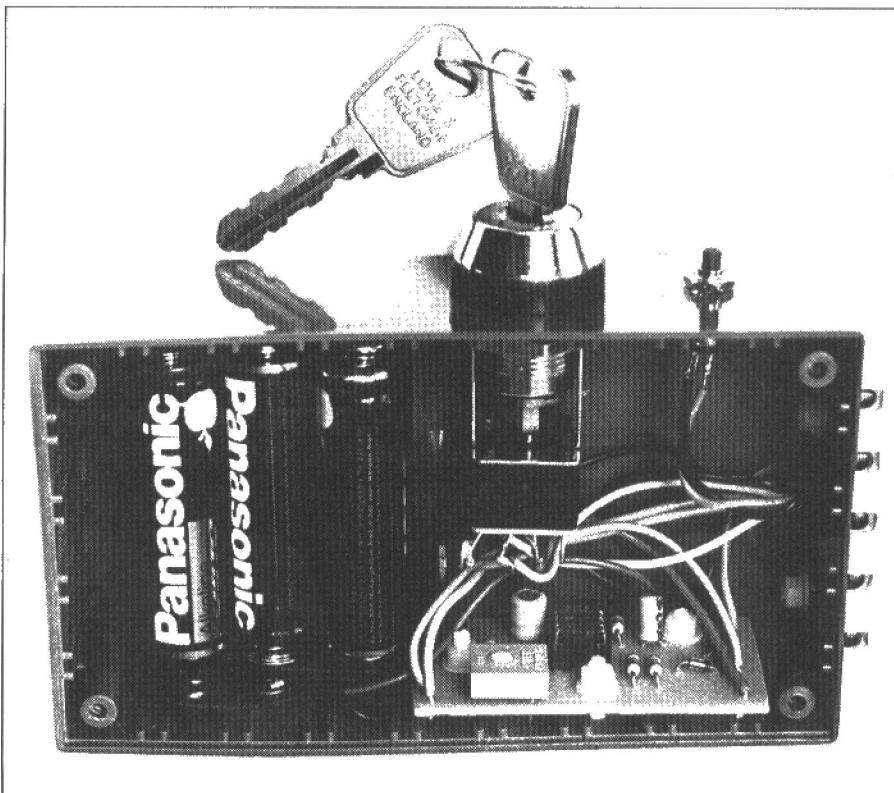
must have sufficient current rating to handle the fog lights - see Parts List.

To be effective, the sounder must be very loud and have a low current requirement. The device specified in the parts list worked very well in the prototype unit. Beware of small cheap buzzers which will not be loud enough for the purpose.

Construction

The printed circuit board topside component layout and full-size underside copper foil master pattern are shown in Figure 2. This board is available from EPE PCB Service, code XXX.

Begin by drilling the three mounting holes. Now, hold the panel in position outside the box (component side down) and mark through the holes which will be used to mount it later. This is difficult to



do from the inside, especially when the components are in position.

Solder the IC socket in place. Follow this with all resistors and capacitors, then solder the diode and transistor, taking care over the orientation. The transistor is amply-rated and does not need a heat sink. Note that C2 and C3 are electrolytic capacitors and must be connected with the correct polarity.

Drilling holes

Solder 10cm pieces of light-duty stranded connecting wire to the pads

marked S1a, S1b and S2 (2 off). Solder the battery snap negative wire to the pad marked Batt -. Solder similar wires to the points labelled TB1/3, TB1/4 and TB1/5. Drill a hole in the box for the external wires to pass through to TB1. Drill the holes already marked for circuit panel mounting and secure it using small fixings. Drill holes and mount S1, S2 and TB1 (see photograph). Refer to Figure 3 and complete the interwiring. Note that the connections between S1 pole a and TB1/1 and TB1/2 must be made with stranded wire of 5A rating minimum.

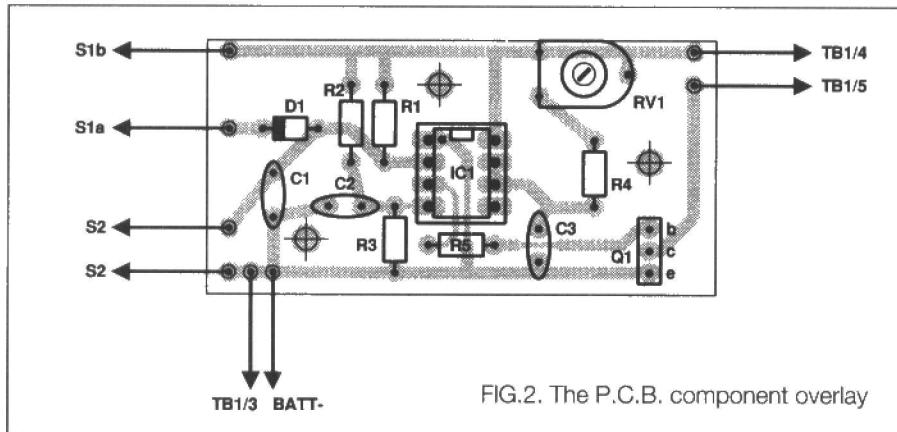


FIG.2. The P.C.B. component overlay

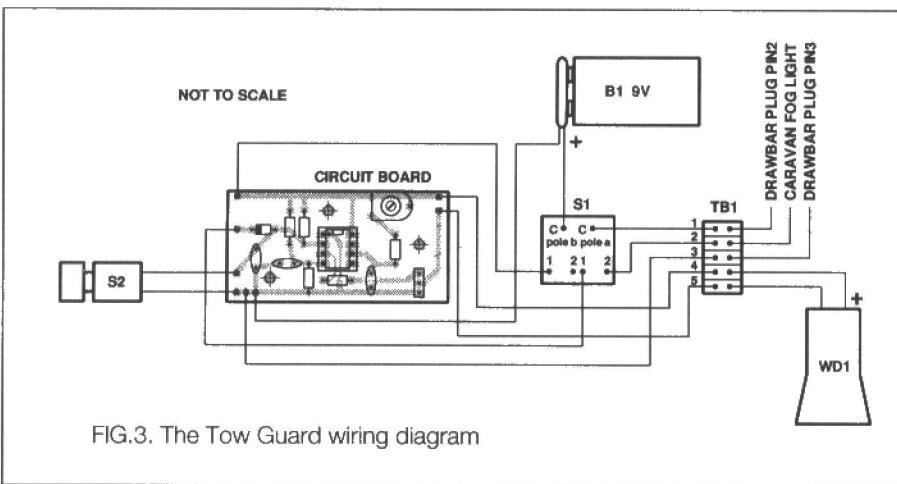


FIG.3. The Tow Guard wiring diagram

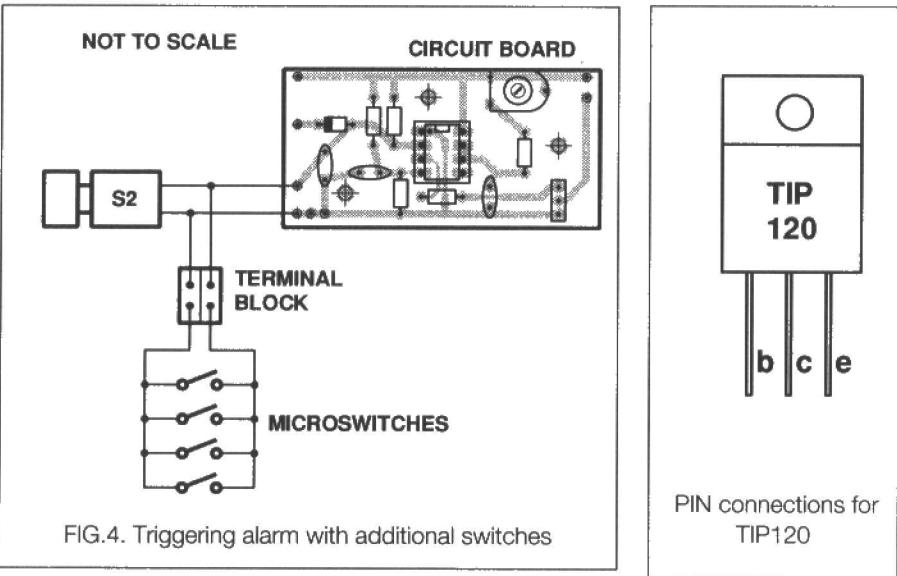


FIG.4. Triggering alarm with additional switches

PARTS LI

Resistors

| | |
|----|------|
| R1 | 10M |
| R2 | 1M |
| R3 | 220k |
| R4 | 470k |
| R5 | 10k |

Capacitors

C1 100n ceramic
C2 4m7 35V pcb elect.
C3 47m 16V pcb elect.

Semiconductors

D1 1N4148 signal diode
Q1 TIP120 (or 121 or 122)
npn Darlington transistor
IC1 ICM7555 CMOS timer

Misceeeaneous

S1 DPDT key-operated switch. 5A d.c. rating - see text

S2 Miniature push-to-make switch

B1 Battery holder for 6 off AA alkaline cells and cells to fit.

WD1 Piezo sounder 6 - 12V 300mA operation. 115dB output at 1m.

TB1 5A screw terminal block

8-pin dIL socket, PCB materials, stranded connecting wire, 5A auto-type wire, bullet connectors, Scotchlocks, 12V bulb for testing, solder, etc. Plastic box size 125 x 70 x 49 mm.

This is because it must be capable of carrying the current for two fog light bulbs.

Adjust VR1 fully clockwise - this will provide minimum timing which is most suitable for testing purposes. Insert IC1 into its socket observing the orientation. Do this without touching the pins because it is a CMOS device and could be damaged by static charge which may exist on the body. Alternatively, touch something which is earthed first. With the specified box, the battery holder needs no support since it is a tight push fit. However, if a different enclosure is used, it could be secured using Velcro fixing pads.

Decide on the best place for the main unit. A good plan would be to make a hole in the front of a cupboard and mount it behind this so that only the switches are visible. It must not be far from some point on the wire running

underneath the caravan which feeds the fog lights.

At the same time, decide how and where the sounder should be fitted. In the prototype, it was mounted in a sheltered position under the caravan. Where wires are run under the caravan, they must be of light-duty automotive type and secured with small cable clips.

The right connections

Locate the fog light feed wire leading from pin 2 of the 12N plug - it will be coloured blue. Find also the earth (chassis) wire leading from pin 3. This will be white. The blue wire will need to be broken or disconnected at some convenient point and auto-type 'bullet' connectors used to extend the free ends to reach TB1/1 and TB1/2 through a hole in the floor. Use wire of a similar type to that used for the existing fog light wiring. The end of the wire leading to the plug should be connected to TB1/1 and that leading to the fog lights, to TB1/2. A connection will need to be made to the white (earth) wire, without breaking it and taken to TB1/3. This could be done at an existing terminal block or a Scotchlok connector could be used. Scotchlocks can be bought from any car accessory shop.

Alarming experience

For testing purposes, it is strongly advised to use a small (2.2W) 12V bulb in a suitable lamp holder in place of the siren. You will thus maintain your friendship with the neighbours!

Set S1 to off. Insert 6 alkaline AA size cells into the battery holder, connect the battery snap and switch on. If the alarm happens to self-trigger, switch off and wait for a few seconds before switching on again. Make a basic test by pressing the test button - the lamp should light for around 15 to 30 seconds then go off. Test at maximum timing - this should be between 3 and 6 minutes. Adjust VR1 for the timing required.

If windows, gas bottle cover, etc., need to be protected, this may be done using lever-arm microswitches. A pair of contacts is used, and they close when the protected item is disturbed. These are connected in parallel with test switch, S2 (see Figure 4). Two extra sections of terminal block could be used to make the connections. When everything is working satisfactorily, connect the sounder, observing the polarity - TB1/4 to the positive (red) wire and TB1/5 to the negative (black one) extending the wires as necessary using bullet connectors. Check that the unit triggers when the plug is slightly

inserted into the tow bar socket.

There are two conditions where the alarm will fail to trigger. Firstly, if both bulbs in the tow-car's fog light system have blown (since one of these is needed to make IC1 pin 2 go low) and, secondly, if the fog light is operating when the plug is inserted. This would give a high state which would be blocked by reverse-biased diode D1 and have no effect. These are both thought to be rather unlikely circumstances.

Since water conducts electricity to some extent, the caravan plug was sprayed with water to check that this did not trigger the alarm. It did not, although it is thought that this could possibly happen. The plug should therefore always be protected from rain. Wise caravanners will do this anyway to prevent corrosion and the need to replace the plug after a short time.

It would be a good idea to make a bold WARNING: ALARM ARMED sign for the caravan hitch. This will remind you to switch off the unit and might even deter the thief from proceeding further! Finally, don't forget to test the alarm every few weeks. The moment there is any sign of weakening sound, the batteries should be replaced.



THE IDEAL TOWN HOPPER

No, it's not a Sinclair C5, and yes it's electric, environmentally sound and easily manoeuvrable, the City-e! explodes the myth that electric cars are impractical, with a 50KM range and 50KM top speed, it's the best alternative in urban traffic, 10 metres of flex and conventional plug ensure that wherever you are, just plug in.

- ★ £2,499 CABRIO OR HARDTOP
- ★ £1,875 EX-DEMONSTRATORS CABRIO-HARDTOP
- ★ VEHICLE TAX EXEMPT
- ★ CHARGING TIME 100% 8 HOURS 75% 3 HOURS
- ★ CHARGING COST 9 PENCE/100% CHARGE



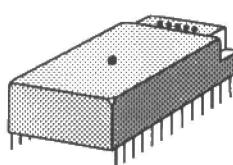
Patriotic Street, St. Helier, Jersey, Channel Islands
Telephone: 0534 80973 Facsimile: 0534 80974

PROGRAMMER EPROM, EEPROM & FLASH for all types 2k (2716) to 8 meg (27C080) 12.5 to 25 volt. Connects to printer port on IBM, fast program (275.2 ± 45 secs). Comprehensive editor software. Requires 15-30 volt AC or DC power supply (optional extra £6.50) **only £99.95**
ROM EMULATOR. Works on ANY computer with a printer port. Code sent to printer appears in rom on target board. Fast download, switchable to emulate up to 32 ROMs. **only £39.95**
PC SCOPE. Fast 8 bit A/D converter connects to printer port on IBM PC. Supplied with Scope and Voltmeter software with Data and A/D routines to write your own programs. **only £29.95**
PIC16C54/55/56/57/84/71 PROGRAMMER. Runs on IBM PC, with editor software, protection fuse blowing etc. requires 15-20 volt AC or DC power supply (optional extra £6.50) **only £69.95**
LOW COST PIC16C54/56/71/84 I.C.E. Connects to printer port on IBM PC. In circuit emulator with single step; go with breakpoints etc. Supplied with connection information and PIC DEV 54 or PIC DEV 71 software as below. **only £89.95**
PIC DEV 54. Integrated Text editor/Assembler/Disassembler/Simulator for 54/55/56/57/87 Pic's. Runs on IBM PC. Single step trace, go with breakpoint etc. etc. Everything you need to write and debug programs. Runs with or without the I.C.E. unit. **Software only £29.95**
PIC DEV 71. As above but for PIC16C71/84 controllers. **Software only £29.95**
ASM51 + SIM51. Text Editor/Assembler plus simulator for Micro controllers 8051/52. Runs on IBM PC. Full control of registers & memory. Label disassembly, program trace, single step etc. **Software only £19.95**
ASM48 + SIM48. Same as Asm51 but for 8748/49 Micro controllers. **Software only £19.95**
All hardware carries a 12 month parts & labour Guarantee. S.A.E. or telephone for further details.
Please add £1.50 for Post/ Pacing.

JOHN MORRISON (MICROS) Dept ETI.
4 REIN GARDENS, TINGLEY, WEST YORKSHIRE WF3 1JR
Telephone 0532 537 507



EPROM EMULATOR



Only 11mm high, it plugs directly into your Eprom socket. Compatible with standard Eproms. Allows faster product development.

Data is written to the emulator directly from your computer, via a 4mm cable which plugs into the emulator.

The MicroRom Eprom Emulator offers non-volatile storage and can download a 27256 file in only 1 SECOND. Prices from £99 + VAT

TEL / FAX 081 880 9889

SquareWave Electronics Ltd., Imperial House, 64 Willoughby Lane, LONDON N17 0SP

MAGNETISM AND MAGNETOMETERS

Magnetism is a subtle and mysterious force that influences a lot of things around us. Keith Garwell continues his practical exploration of how to measure minute changes in a magnetic field with the construction of a Fluxgate Magnetometer.

The term Fluxgate Magnetometer sounds rather like something from science fiction, but anyone familiar with magnetic amplifiers, or mag-amps as they are often referred to, will also find the workings of the fluxgate magnetometer, or FGM as we shall henceforth call it, familiar. If you are not familiar with such devices, then I shall try and explain the principles involved, but short of writing a whole textbook on the subject we will only be looking at the basics, so there is plenty of room for further experiment and enhancement.

This version of a fluxgate magnetometer has also been designed to be adjustable in gain and offset, so that it can be a bit of a jack of all trades and, with adjustment and enhancement, at least master of some.

For example, its sensitivity is such that it gives about 3V for the horizontal component of the earth's flux, with the design

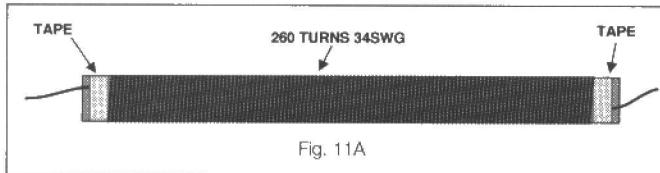


Fig. 11A

sensitivity reduced by 10. If required, therefore, the gain can be altered so that it gives 5V for 50 micro Teslas, i.e. $10\mu\text{T}$ per volt. This is a convenient scale where direct readings of field strength are required, enabling both horizontal and vertical components to be measured.

Similarly, its circular sensitivity is about 3V for 7 degrees as developed (7mV/minute). This sensitivity to rotation means that the normal small changes in the earth's field (5 to 10 mins of arc) would produce an output change of 35 to 70mV.

General Outline of Operation

Earlier, the FGM was described as having 3 windings, two on a

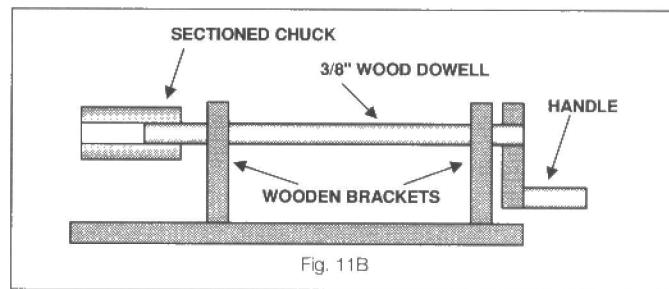


Fig. 11B

Fig. 10 Magnetometer block diagram

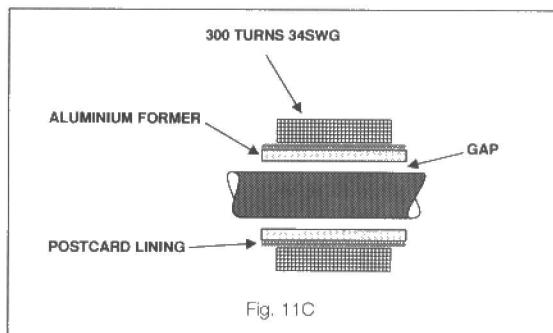
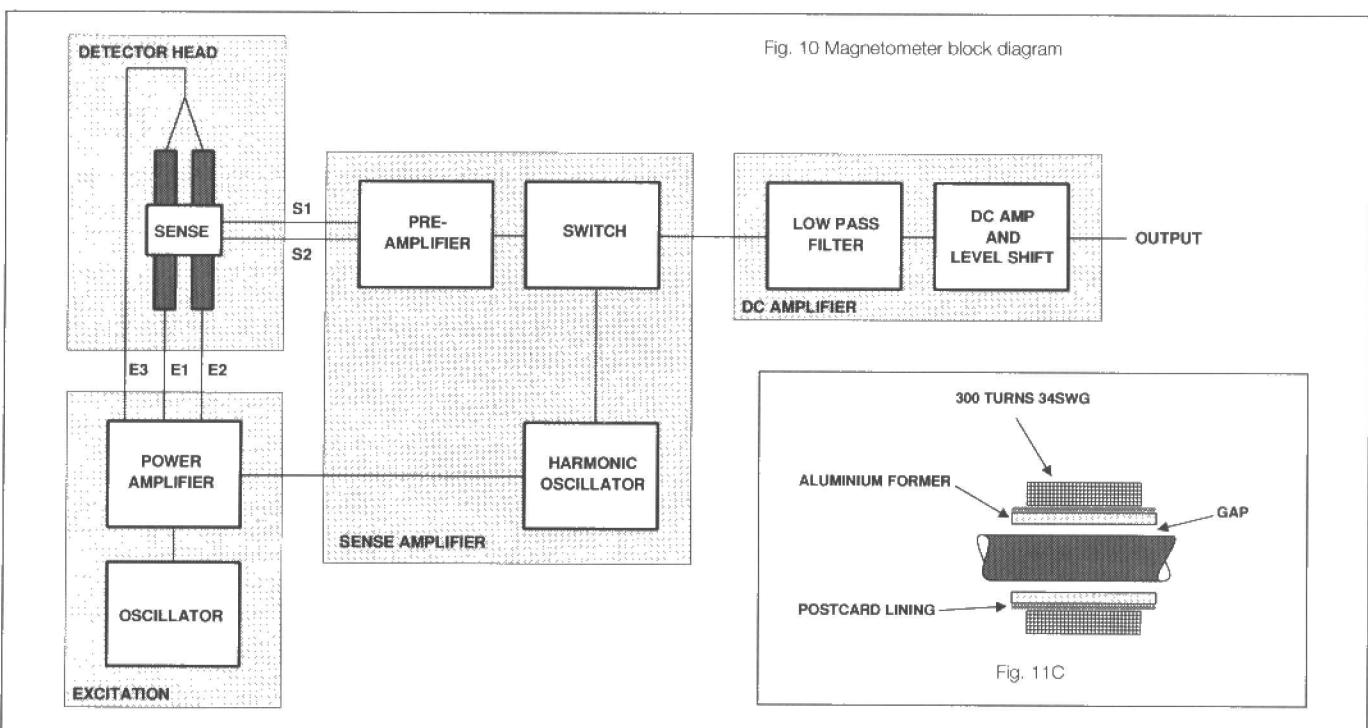


Fig. 11C

pair of ferrite rods (more or less standard size) with these in turn surrounded by a third winding, which is the sense winding. Figure 9a gives a plan view of the arrangement, showing the sense winding surrounding the middle portion of the two rods. The windings on the rods run their whole length and are wound in the same direction on each rod. Figure 9b shows an end on view.

Consequently, if the two ends of the windings are connected together at one end and the excitation applied at the other, the field generated by the rods will be equal and opposite. If this is the case, no voltage would be induced in the sense winding.

If the current through the sense winding is such that the flux generated in the rods just reaches saturation, any external field

would unbalance the system and produce the second harmonic of the excitation. It is the magnitude of the second harmonic which is used to indicate the magnitude of the external field.

In this imperfect world, the two rods and their windings will not be absolutely identical so there will always be some of the fundamental present in the sense winding. Consequently, the measuring system has to include a device for separating second harmonic from fundamental and this is done by means of a switch which inverts the waveform every half cycle at the frequency of the second harmonic.

Figure 10 shows the various parts of the magnetometer enclosed within dotted lines. The arrangement of ferrite rods and the windings we will call the detector or head, top left. This is best in its own enclosure, separate from the rest of the electronics. In fact, five wires are necessary to connect the head to the electronics as the centre connection of the two rods is brought out for balancing, as we shall see later.

The windings on the ferrite rods must be driven with a fairly immaculate sine wave - life is difficult enough without having any second harmonic to cloud the issue. The frequency is not very important, except that since the exciter windings have inductance, the higher the frequency, the higher the AC voltage will have to be to generate the necessary current. This starts to become a problem, particularly if the equipment is to be portable. It's as well not to have it anywhere near 50H or a multiple, otherwise it will detect any AC mains fields in the locality. As we will see, using common component values gives 71H as the exciter frequency, which is very convenient.

Figure 10 shows that this is achieved by using a separate sine wave oscillator, followed by a power amplifier. About 300 mA AC is required in the excitation windings and about 3 V AC is required to generate this current, i.e. a bit less than a watt. As we shall see, there is a convenient chip that will do this without requiring many external components and without spoiling the waveform from the oscillator.

Continuing with Figure 10, the sense winding is followed by a

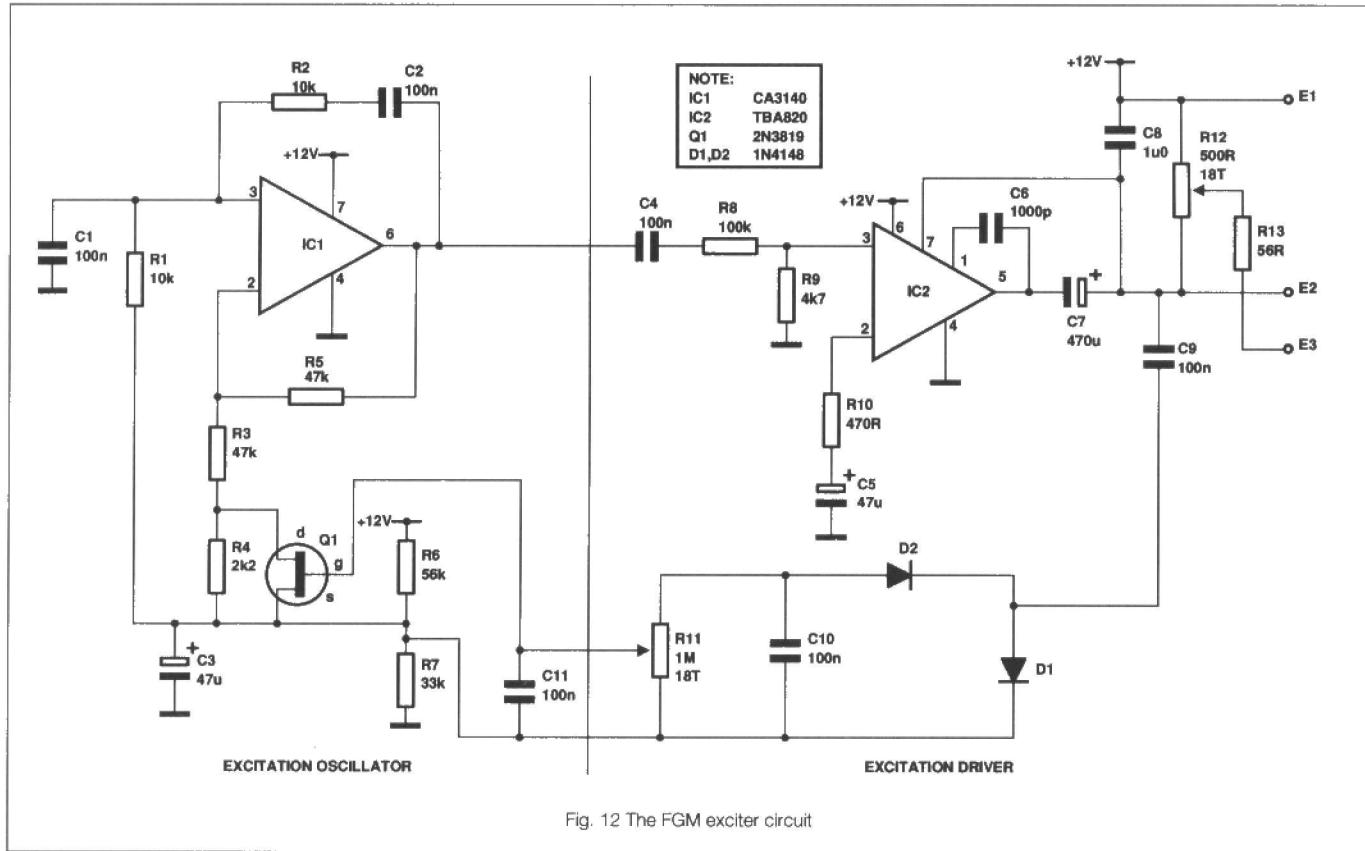
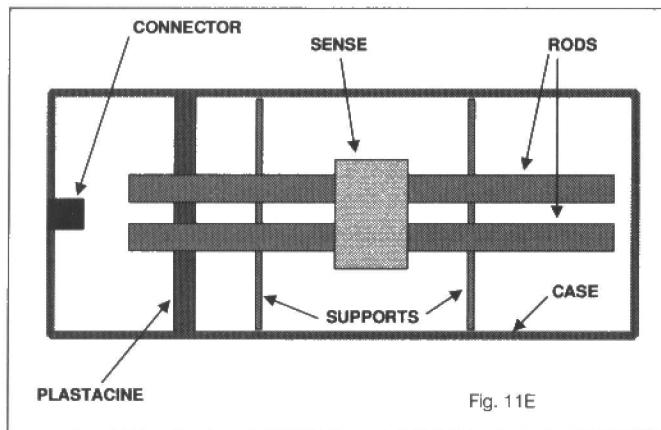
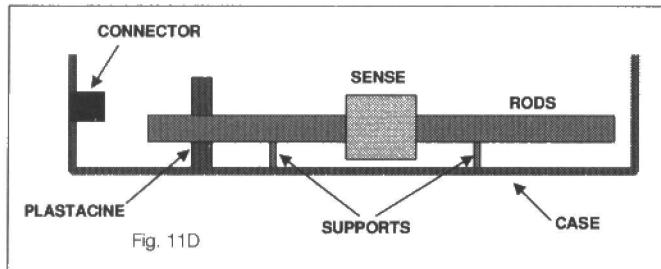
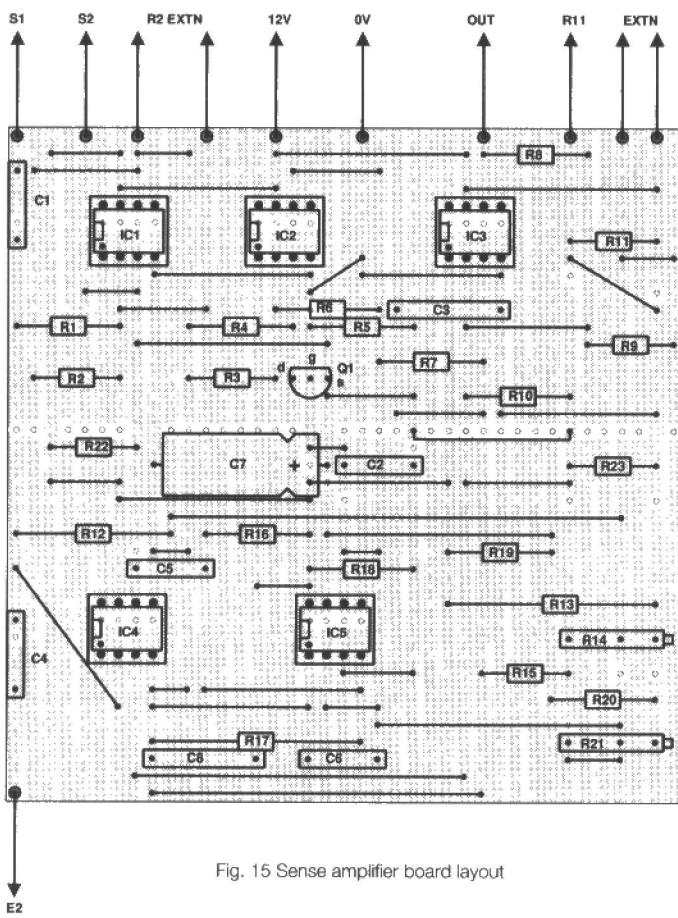
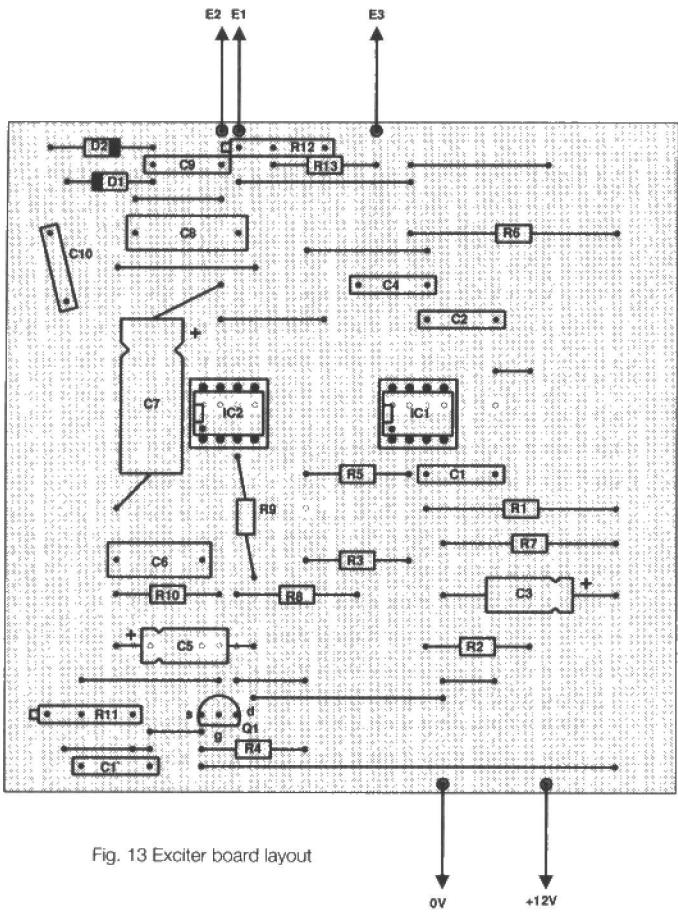


Fig. 12 The FGM exciter circuit



pre-amplifier, with a gain of 100 to get the signal up to a reasonable size before measurement.

Now one of the more crafty bits. The output of the sense amp is fed into an amplifier whose gain can be switched between +1 and -1. If the switching is at the same frequency and in the correct phase with the second harmonic, the harmonic will appear at the output of the switch as DC, anything else remaining as AC.

The two are easily separated by following the switch with a low pass filter. To enable the output from the filter to be referenced to the 0V or ground rail and also to allow the level to be shifted (offset) to suit the measurements in progress, the filter is followed by a DC amplifier.

The final part and arguably the most important, is the oscillator which drives the switch. This must be stable and remain at the second harmonic, while at the same time being capable of synchronisation with the output from the power amp which is providing the excitation. Our old friend the 555 comes to the rescue.

Making the FGM

The Detector

The point must be made that during the construction of the detector head, the end objective must be a mechanically rigid structure. Remember, we are considering a device capable of detecting minutes of arc of movement. My advice is that this can only be achieved easily by constructing the device so that it can be potted in resin when finished and the following constructional details assume that this will be the case.

I would also suggest that the potting is not done until the rest of the device is complete, so that it can all be tested before being set in concrete as it were.

The first requirement is, of course, a pair of ferrite rods. These should be about 100mm long and 8mm in diameter. Since they are only about 50 pence each, it is worth getting three and then selecting the two that look as near identical as possible.

Figure 11a shows the arrangement, which is quite simple. Using 34 gauge wire, anchor one end with tape. Close wound, 260 turns of 34 gauge takes up about 64mm, leaving room at either end for the tape. I used 3/4in masking tape.

Close winding this length is a bit tedious if you don't own a winding machine, which I don't, although I do have a large hand drill with a half-inch chuck. If you have one of these, the technique is to clamp it to a work table/bench. Then, with the wire taped to one end, this taped end can be inserted into the chuck. The tape also helps to make the grip firm, as the ferrite and the chuck jaws are so hard. Having counted the number of teeth on the hand wheel and chuck wheel, 260 x chuck wheel teeth/hand wheel teeth will give the number of turns required. A little bit of experimenting enables a suitable technique to be found to allow the wire to be laid nicely on the ferrite. Try to make the two rods look as similar as possible - they will be electrically trimmed by the electronics later.

Failing a suitable drill, it is still worthwhile to try and bodge something up so that the rod can be turned

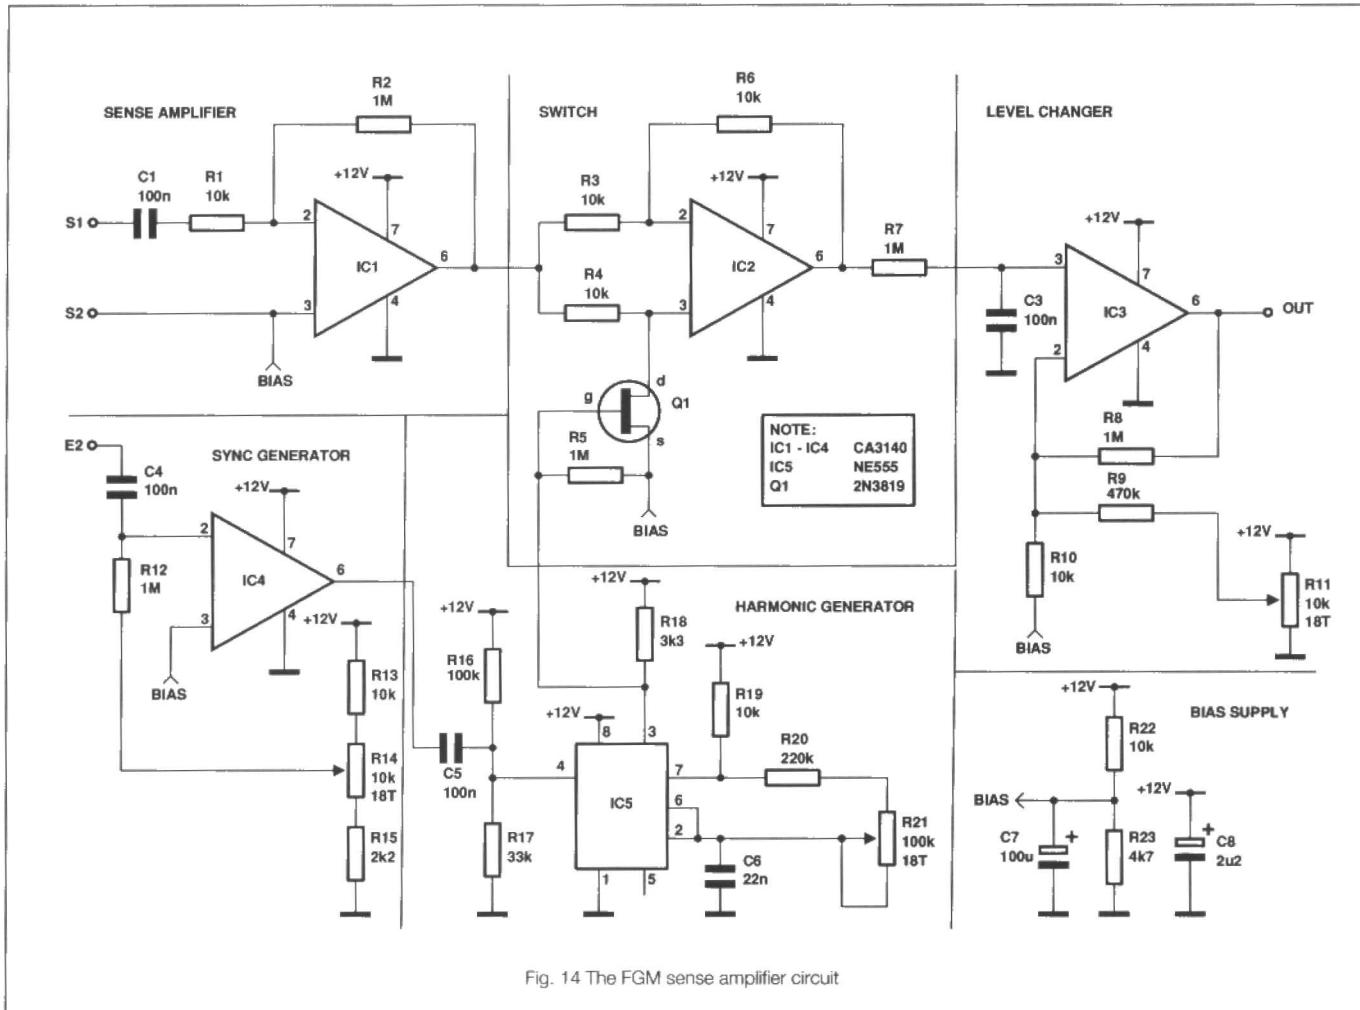


Fig. 14 The FGM sense amplifier circuit

comfortably. I tried winding by hand initially and found it very difficult to keep the turns tight and close together. Figure 11b contains a useful suggestion, in which a length of 3/8in wooden dowel turns in two wooden brackets. A handle at one end and a chuck consisting of a piece of broom-handle with a 3/8in hole through it at the other, completes the general idea. Alternatively it might be worth while making something of the same sort from Meccano or Lego. If the device is not sufficiently rigid, perhaps adding a steady to the ferrite rod at the end remote from the chuck may help.

I should now let the cat out of the bag and be honest. I bought the rods, already had some 34 gauge, fastened wire and rod in the chuck and then wound, while counting the turns of the handle, until I reached the other end with enough room to fasten off. I worked out the number of turns on the rods afterwards! However the rest of the design is dependant on the rods, so don't

depart from the size and numbers of turns, unless you are prepared to change the odd component here and there.

I will say that I am not aware of any reason why only 34 gauge should be used, provided that the number of turns is correct and the winding is spread evenly over the rods.

The next thing required is a former, onto which the wire for the sense winding can be wound. This I made by folding some aluminium sheet so that I ended up with the item shown in Figure 11c. This is made just large enough to accommodate the two rods and the width for the winding is 13mm (1/2in). When building it, I felt that using a metal former would be helpful if any screening was needed between the excitation and sense windings.

The other side of the coin is that aluminium sheet is as easy

to use as anything else and is wholly self supporting - no glue needed. However, there is one most important point. There must be an insulating gap between the two ends, otherwise the former will be a short circuit and will screen the sense winding from everything!

I made the former by first cutting and folding a strip of sheet so as to make a length of U section and then marked off the points that would become the corners. The channel had been made deliberately longer than required by the four sides, so that a short flap could be left. Using the tin snips to cut through the sides at the marked points allows the channel to be bent into a frame, with the little flap left tucked under the start.

The next move is to use an ordinary post-card to provide a lining, made in exactly the same way as the aluminium former but of course very slightly narrower so that it will sit inside the aluminium. Leave a tab end so that it can be tucked into the gap between the ends of the aluminium. This arrangement will ensure that no sharp corner on the aluminium will damage the insulation of the wire. Having fitted the card lining, put on 300 turns of 34 gauge. These can be put on by hand as there is no need to try for perfect layers, so long as it's reasonably tidy.

Having done all three windings, the next stage is to dunk them in hot beeswax. This will stop any movement in the windings. If you have never done this sort of thing before, be careful not to get it on your fingers. It's easiest to heat the wax in a saucepan over an electric plate, but any other metal container will do, so long as its diameter is greater than the length of the rods and not so deep that dipping the windings is awkward.

The wax should be heated until it can be smelt, but it need not be smoking hot. Just hold a rod by its wire ends and lower

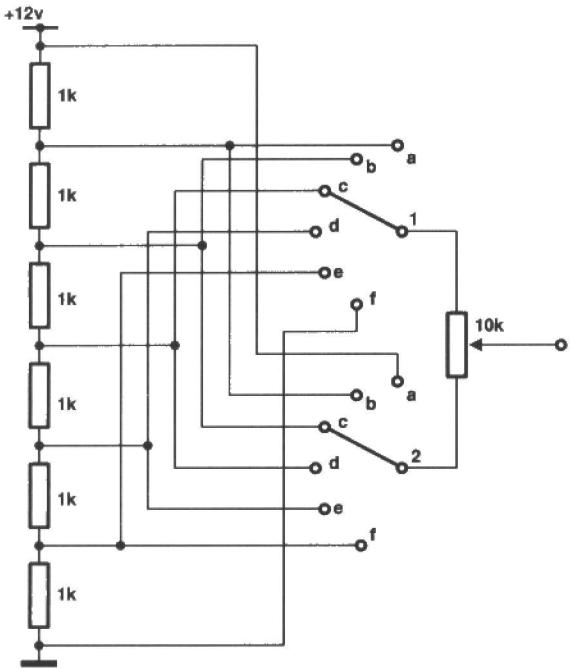


Fig. 17 Alternative to ten turn helipot

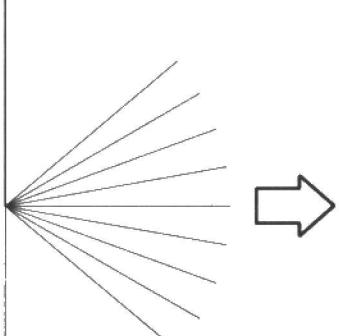


Fig. 18 Compass guide

replaced by wax. Keep it in until the frothing stops, then take it out and stand to cool or suspend it by its wires to cool. It occurs to me in hindsight that all three can be hung to cool by their wires.

The method of assembling the whole thing into a case is suggested in Figure 11d and 11e. I used one of the readily available plastic boxes which have slots on the inside intended for PCBs.

A five pin din socket fitted into one end of the box is the easiest means of connection. However, one which has a steel jacket won't do. Many of them have a moulded jacket which is some sort of alloy and has a finish which is chromium in appearance. These would be fine.

Slide two pieces of plastic, about 1cm wide and as long as the inside width of the box, down the slots so that they will act as supports to the two rods. The sense winding should already be in place on the rods.

As the whole thing is eventually going to be potted in resin, a barrier has to be erected to stop the resin getting into the din

connector. This is most easily done by erecting a plasticine barrier between support and connector (Figure 11d). Arrange the plasticine so that it is higher than the supports and then gently press the rods into position. Add more plasticine above the rods to continue the barrier. Finally, make the connections between the windings and the din connector and make a note of them. There should be five altogether - the two ends of the sense winding, the two ends and the centre tap of the exciter winding.

The FGM Exciter

How it Works

Figure 12 shows the exciter in two parts, the oscillator and the driver. The oscillator is a conventional RC op-amp oscillator where C1, R1, C2 and R2 are the components defining the frequency of oscillation. C1 and C2 are equal, as are R1 and R2. As we are not particular to the precise frequency, common values are used C1 = 0.1µF and R1 = 22K.

There are two slight complications to the normal circuit, the bias and the level control. Resistors R6 and R7 set the bias level midway between the limits of the op-amps output 0V and 9V. These limits are set by the op-amp used, a CA3140 which while it will go down to the lower supply level will only go up to within 2 to 3V of the upper level. C3 provides decoupling to this bias supply.

Whether IC1 will oscillate or not depends upon the circuit gain. If it's greater than 1, then it oscillates. The gain of the circuit is set by R3, R4 and R5. With all three in circuit, the gain is less than one and it won't oscillate. If R4 is shorted out the gain is higher than 1 and, low and behold, it does oscillate.

R4 is not shorted out but its value is modified by the shunt value of FET1 which is modified by the driver output, as we shall see in a moment. IC1 never reaches its output limits, so the waveform is never distorted by any clipping.

On to the driver. This is a nearly conventional circuit using the TBA820 which is an audio driver. Its attraction is its low distortion and the way in which its internal safeguards work. The latter trigger at the crucial overload levels, so that cut-off is sharp, whereas with some audio ICs the cut-off is progressive. As the IC works fairly close to its maximum dissipation, we cannot afford to have the 'Hi-Fi' spoilt by progressive shutdown.

R8 and R9 set the input level and R10 and C5 set the IC's gain. C6 is a compensation capacitor, its value being rather higher than normal to accommodate the low frequency being handled. C7 provides DC isolation and a bootstrap facility to pin 7.

C8 is much larger than usual, on account of the highly inductive load - the two excitation windings on the ferrite rods which are connected to E1 and E2, the centre connection goes to E3.

The actual output of the driver has to be maintained within 1 part in 200. This is achieved by the self regulation as follows.

C9 takes off the voltage appearing across the rod windings and it is rectified to provide a negative voltage across R11 by D1, D2 and C10. The DC voltage at R11 is used to modify the gain in the oscillator via the FET.

Finally, although not strictly part of the driver, R12 and R13 are included for convenience and are used when setting up to balance the excitation between the two ferrite rods. This will be dealt with later under setting up, as will the adjustment of R11.

Construction

Figure 13 shows the layout using 39 x 39 matrix board. I must admit I'm not madly enthusiastic about matrix board, but for a one off or for something which is possibly subject to change I

haven't found anything better. For me, it's very difficult to lay out, but there's madly special about this one. It's all low frequency, so no problems there.

There are only 15 tracks to cut - four between the pins of each IC, one just to the right of IC1 below the link, one between the left hand ends of R3 and R5 and then a little collection of five more, shown as being behind C5.

The outlines do not indicate the shape of the components, they indicate the connection points, so R6 is not that big, but the ends of the rectangle show the tracks to which it is connected. Having said this, there are two exceptions, the capacitors. C7 is arranged vertically as a convenience (well have you ever seen a horizontal convenience?) but connected to tracks 6 and 12 as shown. Similarly, C10 is shown vertically but is actually connected across tracks 2 and 3 so is at a slight angle.

Down by C11 there is a little cluster of 3 links, don't miss any.

Testing will be considered separately when we have looked at the next unit - the sense amplifier.

PARTS LIST

Resistors

| | |
|--------|------------|
| R1, R2 | 10K |
| R3 | 47K |
| R4 | 2K2 |
| R5 | 47K |
| R6 | 56K |
| R7 | 33K |
| R8 | 100K |
| R9 | 4K7 |
| R10 | 470R |
| R11 | 1M - 18t |
| R12 | 500R - 18t |
| R13 | 56R |

Capacitors

| | |
|-----|----------|
| C1 | 0.1µF |
| C2 | 0.1µF |
| C3 | 47µF E |
| C4 | 0.1µF |
| C5 | 47µF E |
| C6 | 1000pF P |
| C7 | 470µF E |
| C8 | 1.0µF |
| C9 | 0.1µF |
| C10 | 0.1µF |
| C11 | 0.1µF |

Semiconductors

| | |
|------|--------|
| FET1 | 2N3819 |
| IC1 | CA3140 |
| IC2 | TBA820 |

Diodes

| | |
|--------|--------|
| D1, D2 | 1N4148 |
|--------|--------|

Notes

Resistors suffixed 18t are 18 turn Cermet presets. All others are hi-stabs with no significant power requirement. Maplin 0.6W MF are fine. Capacitors suffixed E are axial electrolytic, C7 must be 25V, others can be 10V. Capacitors suffixed P are poly-styrene. All others are metallised polyester film. All are available from Maplin.

The FGM Sense Amplifier

How it works

Figure 14 shows the electronics which I have called collectively the sense amplifier. In fact, it consists of several sub sections, only one of which is a simple amplifier and indeed I think it best if I start the description with the harmonic generator, bottom centre.

It's the standard arrangement of a 555 wired as an oscillator. C6 is the C which establishes the frequency and R19, 20 and 21 being the corresponding R. The variable R21 enables the frequency to be adjusted to the second harmonic of the the excitation oscillator. R19 is small compared to R20 so that the mark-space ratio of the output is close to 1 to 1.

The reset input at pin 4 is used to achieve synchronisation between the oscillator and the excitation. This is done by the sync generator (bottom left) being connected to the excitation, E2. The sine wave is applied to the inverting input pin 2, together with the DC level established by the resistor chain R13, 14 and 15. The non-inverting input is connected to a bias supply (bottom right) of about 4V. A transition occurs as the sine input becomes greater or less than the bias. Adjusting R14 therefore enables the transition point to be adjusted.

The square wave output is differentiated by C5 in conjunction with R16 and 17. The negative going pulse resets IC5, thus ensuring that it is always in phase with the excitation.

Now back to the beginning, top left. The output from the sense winding is applied to IC1, which is a straightforward inverting amplifier with a gain of 100. The bias avoids the necessity for double power rails, its voltage being approximately half the maximum excursion of the output of IC1 (0V to about 9V).

The output of IC1 is about 6V peak to peak and is applied to the switch IC2. The gain of this stage is switched between +1 and -1 by FET1, driven between the conducting and non-conducting state by the square wave output of the harmonic generator. The presence of the second harmonic in the output of IC1 thus appears as a DC component at the output of the switch.

R7 and C3 form a low pass filter which removes much of the AC components present in the output of the switch. This is an important point which I shall deal with later in much more detail.

The DC output of IC2 is referenced to the bias supply. To enable this reference to be changed and to allow the scale to be expanded, IC3 is a DC amplifier with an adjustable reference, obtained by means of R11.

Construction

Figure 15 shows the layout, again using a piece of 39 x 39 matrix board. The top half (roughly) contains the sense amp, switch, level changer and level adjust pot, while the bottom contains the sync generator, second harmonic oscillator, frequency and sync adjust pots. I am the first to admit that the board is a bit cramped and anyone who opts to construct it on two boards is probably more sensible than I am.

I found the chief difficulty was the necessity for cutting tracks and freely admit to spending an hour or two trying to work out why the sync generator wouldn't work. In due course I discovered that the input capacitor C4 was shorted out by the track. Cutting it between the two connections of C4 worked wonders!

Although Figure 15 shows the track cuts with Xs, I have included the Board layout in the Foils section which shows just the tracking and cuts. Note that this is from the component side.

It is perhaps best to use these two figures as guides and not absolute statements of fact. The best way to build this board is to start at the top, leave a couple of rows blank for connection

to the outside world, i.e. S1, S2 12V, etc., and then work steadily downwards, checking as you go that what you are doing corresponds with the schematic. As you come to a track cut, stick a piece of wire through the hole, turn the board over and cut the track. It seems to be fairly easy to keep the vertical numbering correct but more difficult to preserve the horizontal.

As with the exciter board figure, the outlines do not indicate the shape of the components but the connection points. Don't work on this board until you have read to the end of the article.

Resistors

| | |
|-----|----------------------|
| R1 | 10K |
| R2 | 1M |
| R3 | 10K |
| R4 | 10K |
| R5 | 1M |
| R6 | 10K |
| R7 | 1M |
| R8 | 1M |
| R9 | 470K |
| R10 | 10K |
| R11 | 10K - 18t see later. |
| R12 | 1M |
| R13 | 10K |
| R14 | 10K - 18t |
| R15 | 2K2 |
| R16 | 100K |
| R17 | 33K |
| R18 | 3K3 |
| R19 | 10K |
| R20 | 220K |
| R21 | 100K - 18t |
| R22 | 10K |
| R23 | 4K7 |

Capacitors

| | |
|------|---------------|
| C1-5 | 0.1 μ F |
| C6 | 22000pF |
| C7 | 100 μ F E |

Semiconductors

| | |
|-------|--------|
| IC1-4 | CA3140 |
| IC5 | NE555 |

Notes

R2 and R11 may be implemented off-board see later text. R20 possible alternatives required 180K, 200K, 240K. Resistors suffixed 18t are 18 turn Cermet presets, all others are hi-stabs with no significant power requirement. Maplin 0.6W MF are fine. Capacitors suffixed E are axial electrolytic 10V. All others are metallised polyester film.

Final Construction

So far we have looked at the construction of the boards. Now we must consider the finishing touches. How is the magnetometer to be housed and what sort of duty is it to perform? It is also necessary to decide how the output is indicated.

In the first instance I will assume it will be a moving coil voltmeter of 10V FSD as this is by far the easiest instrument for setting up.

There are two types of duty. Either as a general purpose instrument, in which case the user must have easy control. It must be possible to change its behaviour. If the magnetometer is going to be dedicated to one particular task, then it will probably only need to be set up once, in which case the level changer control can be an 18 turn cermet pot as shown. If the instrument is too sensitive, then its gain may be reduced by changing the value of R2. The gain will be in proportion to the value of R2 and to reduce the sensitivity by a factor of ten, reduce R2 to 100K.

On the other hand, if it is to be a general purpose instrument then the gain will need to be adjusted easily, as well as the level. In this case, R2 can be replaced off-board by switched resistors, the switch being on the front panel of the box or whatever contains the electronics. The level control R11 must also be on the front panel and it will have to be a multi-turn pot, otherwise the control will be too coarse. It will be useful to be able to reverse the sense winding connections and also to disconnect it - a two pole changeover is useful here.

I have ended up with three controls - the sense winding switch with the On-Off-On combination (allowing the sense to be reversed or switched off). This is also very convenient for commissioning the instrument), a gain switch setting the sense amp gain to either 10 or 100 (100K or 1M) and a wire-wound 10 turn for the output control.

The next step is to provide a suitable housing, mount the boards and the controls (if any) and wire up.

While on the subject, if R11 is to be on the front panel, the RS/Electromail catalogue shows a ten turn helipot which is suitable (3W wire-wound at around £4 plus VAT and postage). An alternative is the arrangement shown in Figure 17, which employs a 2 pole 6 position switch, half a dozen equal resistors (1K) and a common or garden 270 degree pot. The only shortcoming of this arrangement is that it does have a blind spot at each transition point, the reason being that the voltage across the 1K resistor which has the variable switched to it is slightly less than that across each of the others. This can be overcome by using a chain containing 1K and 220R resistors alternately. The switching is arranged so that the variable is connected across a pair of resistors at each step.

However, this is getting complicated and in any case it's very unlikely that the blind spots will matter.

Next Month

In the next issue, we will look at commissioning and using the Fluxgate Magnetometer.

MIDI EXPLAINED

In Part 4 of this series, Robert Penfold continues to delve into the mysteries of MIDI

In last month's article we considered the way in which MIDI control messages are coded. This month we continue on the theme of MIDI controls, with a detailed discussion of the way in which they are used. The original SCI MIDI specification of 1983 was rather vague about the functions of MIDI controls. It assigned control number 0 to pitch bending, but this function was given its own channel message by the time the first 'proper' MIDI specification was published. No other recommendations were made, which left it up to individual manufacturers to use MIDI controls in practically any way they wished. Some manufacturers used these controls as a means of altering the sound generator settings via MIDI. Others used them for more general control master volume, switching effects on and off, etc.

Some equipment had various functions that could be accessed by control change messages, with any control assignable to any control number. This permitted a large degree of compatibility with other equipment, provided you were prepared to set up the assignable controls to match the other equipment in the system. Gradually, a few conventions came into being and these were largely influenced by the control assignments of the more popular instruments of the time (particularly the Yamaha DX7). Not all instruments had control assignments that kept to these conventions, however.

Strict Control

Eventually, the MIDI 1.0 Detailed Specification was published and this lays down specific uses for many of the control numbers. When using MIDI controls, it is as well to bear in mind that some instruments of the 1980s do not fully conform to the current specification. Instruments manufactured in the 1990s should fully adhere to the recommendations.

Under the current scheme of things, most of the MIDI controls are not used for adjusting the sound generator circuits of a synthesiser. There are a few controls which can be used for this purpose and they have to be used in the manner described in the detailed specification, but this is a subject we will consider in detail later. Most of the MIDI controls are used for more general control, such as master volume, balance, sustain and soft pedals, etc. Table 1 shows the current recommendations for the MIDI control functions. The control numbers that are not included in Table 1 have not yet been assigned a function.

Table 1

| Control Number | Function |
|----------------|--------------------------|
| 1 | Modulation Wheel |
| 2 | Breath Controller |
| 4 | Foot Pedal |
| 5 | Portamento Time |
| 6 | Data Entry Knob |
| 7 | Main Volume |
| 8 | Balance |
| 10 | Pan |
| 11 | Expression Controller |
| 12 | Effect Control 1 |
| 13 | Effect Control 2 |
| 16-19 | General Purpose |
| 32-63 | LSB For Controls 0 to 31 |
| 64 | Sustain Pedal |
| 65 | Portamento |
| 66 | Sostenuto |
| 67 | Soft Pedal |
| 69 | Hold 2 |
| 80-83 | General Purpose |
| 91 | External Effects Depth |
| 92 | Tremolo Depth |
| 93 | Chorus Depth |
| 94 | Celeste (Detune) Depth |
| 95 | Phase Depth |
| 96 | Data Increment |
| 97 | Data Decrement |
| 98 | Non-Registered Parameter |
| 99 | LSB |
| 100 | Non-Registered Parameter |
| 101 | MSB |
| 121-127 | Registered Parameter LSB |
| | Registered Parameter MSB |
| | Channel Mode Messages |

There is a slight problem in assigning specific functions to the MIDI controls. Although these controls are mainly used with instruments for such things as the sustain pedal and controlling external effects units, they are also used to control audio mixers, lighting systems and various other non-musical pieces of equipment. It is still permissible for the MIDI controls to be used for non-musical functions, but the equipment manuals must clearly indicate that they are being used in a non-standard fashion.

Sound Control

Although MIDI controls should no longer be used arbitrarily to control the sound generator circuits of a synthesiser, there are some control numbers specifically set aside for this purpose. There are two types of control, which are the registered parameters and the non-registered parameters. The basic idea is that, wherever possible, a standard set of controls should be used. These are the registered parameters. Unfortunately, with the instrument manufacturers using a variety of sound synthesis methods, there is only very limited scope for a standard set of controls. As yet, there are only five registered parameters - see Table 2.

Table 2

| Registered Parameter No. | Function |
|--------------------------|------------------------|
| 0 | Pitch Bend Sensitivity |
| 1 | Fine Tuning |
| 2 | Coarse Tuning |
| 3 | Change Tuning Program |
| 4 | Change Tuning Bank |

In practice, it is the non-registered parameters that you are more likely to use. These can control virtually anything the equipment manufacturer desires and can be assigned to the parameter numbers in any way the manufacturer sees fit. However, it is a requirement that the equipment manuals must give a full list of parameters and their numbers.

In order to change the value of a registered parameter, first the number of that parameter must be written to controls 100 (LSB) and 101 (MSB). This gives a 14 bit value and up to some 16384 different control functions. The new value for the selected control is written to control number 6 (MSB) and 38 (LSB). Of course, if 7 bit resolution is sufficient, only control 6 needs to be used, but a dummy value should be written to control number 38, otherwise the receiving device might wait for this

LSB before implementing the new value. Obviously, it would fail to implement the new value if no LSB was received.

Much the same system is used to alter a non-registered parameter, but the parameter number is written to controls 98 (LSB) and 99 (MSB). The value for the selected control is then

written to controls 6 and 38, as for a new registered parameter value. There is an alternative method of altering a registered or non-registered parameter. The parameter to be altered is selected in the normal way, but the value is then incremented by the value written to control number 96, or decremented by the value written to control number 97. This method is presumably included as a quick means of making minor changes to the setting of a control.

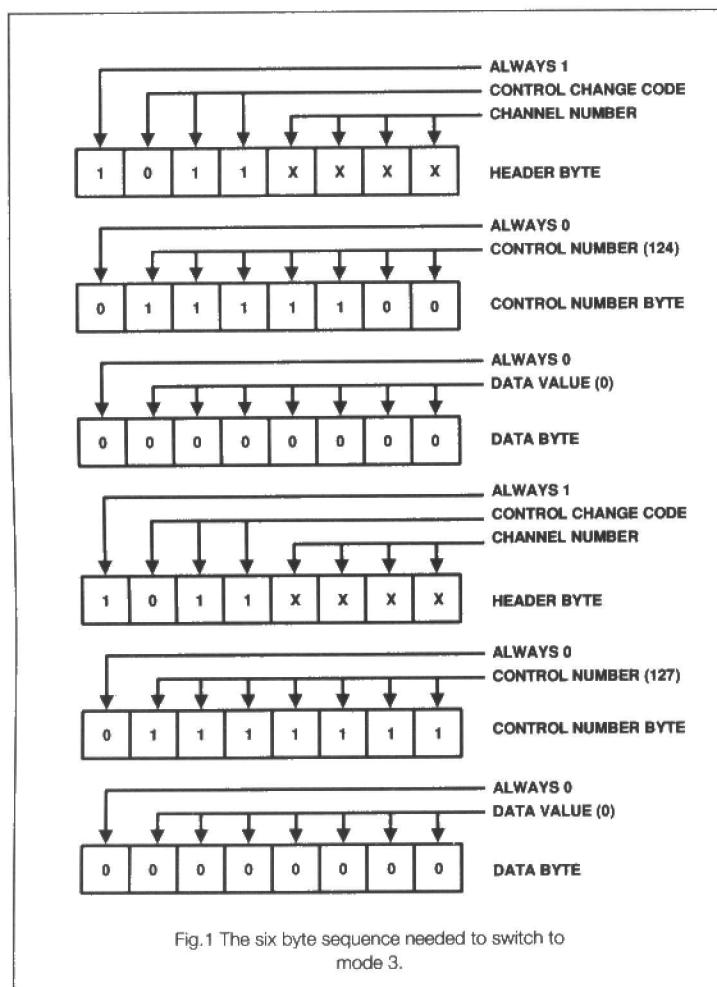
It is acceptable for devices to respond to changes in registered parameter values at start-up, but the reception of non-registered parameter messages should be disabled at power up. The point of this is to avoid accidentally scrambling the sound generator settings of an instrument. When using MIDI, you need to bear in mind that operating the controls of a device in the system will often result in the appropriate MIDI messages being sent. For example, changing an instrument from one program to another will usually result in the appropriate program change message being sent. This sort of thing is generally useful, since it provides an easy means of generating a wide range of messages. It can, however, provide problems for the unwary.

Altering the sound generator settings of an instrument might result in a series of control change messages being sent. If received by the right type of instrument these would presumably result in the slave instrument changing its sound in sympathy with the master instrument.

The effect of the control change messages is unpredictable if the two instruments are of different types, but there is a strong possibility that the sound generator settings of the slave instrument would be well and truly scrambled. If you wish to change non-registered parameters, it will first be necessary to use the appropriate control settings to switch on this feature. It is advisable to switch it off again once you have made the necessary changes.

The degree to which the sound generators of an instrument can be controlled via MIDI control change messages varies substantially from one instrument to another. Some provide quite comprehensive control via this route, but others largely or totally ignore this method. Control via system exclusive messages seems to be the preferred method these days. This may seem

to go against the general spirit of MIDI, which encourages a large degree of standardisation between instruments from different manufacturers. On the other hand, the non-registered parameters are not standardised and it is not practical for them to be standardised to a worthwhile degree. Therefore, regard-



less of which method is used, any large scale alteration of the sound generator settings normally requires some form of dedicated programmer or control program.

Channel Mode Messages

As can be seen from Table 1, the control numbers from 121 to 127 are used for 'channel mode messages'. These are really a separate category of message, rather than being a normal control change message. These messages still use the same three byte format, but in most cases the final byte is always 0 and is really just a dummy byte. The channel mode messages are used for changing mode via MIDI, plus a few other functions. Table 3 gives details of the seven channel mode messages.

Table 3

| Control Number | Function Data |
|----------------|---------------------------------|
| 121 | Reset All Controls Always 0 |
| 122 | Local Control 0 = Off, 127 = On |
| 123 | All Notes Off Always 0 |
| 124 | Omni Mode Off Always 0 |
| 125 | Omni Mode On Always 0 |
| 126 | Mono Mode On Number Of Channels |
| 127 | Poly Mode On Always 0 |

The reset all controls message is a relatively new one, and it simply sets all the MIDI controls back to their default settings. It is presumably included to provide a quick means of resetting the controls back to sensible states if their settings should accidentally become scrambled. The all notes off message is not intended to be used as a normal method of switching off notes. It is a message that can be sent if a malfunction should occur, with notes being left 'droning'.

Controls from 124 to 127 are used for MIDI mode changing, and this operates on the basis of switching 'Omni' on or off, and switching between 'Poly' and 'Mono'. If you require mode 3 for example, there is no mode 3 message as such. Instead you would use controls 124 and 127 to switch Omni off and Poly on. This six byte sequence is shown in Figure 1. When using these messages it is clearly better to think in terms of the new mode names, rather than the old names or mode numbers. Mode change messages terminate any notes that are playing at the time. MIDI instruments do not always respond rapidly to mode change messages, so it is best to allow a second or two after one of these messages before continuing to play an instrument. The third byte in the Mono on message indicates the number of channels to be set to mono mode. In most cases, this byte will have a value of zero and all the available channels are then set to mono mode.

Local Control?

The local control on/off message can also use a value of other than zero in the third byte. This operates just like an ordinary MIDI switch type control, with a value of 0 being used to switch off local control and a value of 127 being used to switch it on. But what exactly is local control? For most instruments it is the keyboard, but not all MIDI instruments are keyboard types. Local control is any built-in means of playing an instrument (the strings and fretboard of a MIDI guitar for example). Here we will only consider local control in terms of a keyboard instrument, but the general principle applies to any form of instrument that has some form of local control.

With local control switched on, the instrument functions

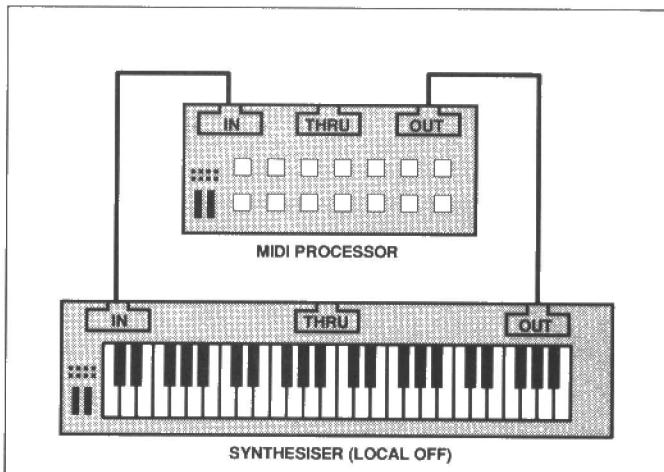


Fig.2 It can be necessary to use local off when using a MIDI processor.

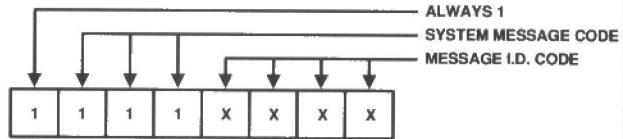


Fig.3 The coding used for a system message header byte.

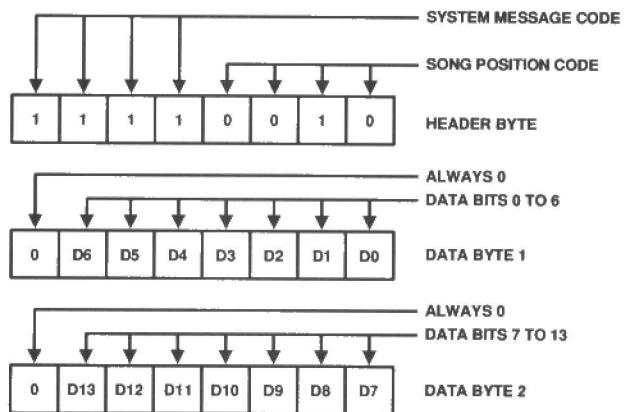


Fig.4 The method of coding used for song position pointer messages.

normally. Playing notes on the keyboard produces the appropriate response from the sound generator circuits and suitable MIDI messages are transmitted from the MIDI OUT socket. With local control switched off, playing notes on the keyboard results in the appropriate MIDI messages being transmitted at the MIDI OUT socket, and the sound generator circuits will respond to messages received at the MIDI IN socket. The sound generators do not respond to the keyboard though. In effect, switching off local control converts the instrument into a separate MIDI keyboard and synthesiser module.

This is perhaps not a feature that most users will require very often, but it does have its uses. When using a synthesiser as a sound module, it is not a bad idea to switch off local control so that there is no risk of unwanted notes being played by someone accidentally leaning on the keyboard. It can also be useful in an arrangement of the type shown in Figure 2. Here, the MIDI output signal is being processed and fed back into synthesiser. If the synthesiser must play the undocored notes and the processed signal, local control should be on. If only the processed notes should be played, local control should be switched off.

System Messages

So far we have only considered channel messages, which are the messages that go through the MIDI system in large numbers during everyday use. The system messages tend to have more specialised functions and it has to be said that in a few cases they are not particularly relevant to current MIDI systems. However, the system messages are certainly a very important aspect of MIDI, and are something that all MIDI users need to understand in reasonable detail.

System messages are directed at every slave unit in the system. Like any other MIDI message, not everything in the system will necessarily respond to them, but everything should process these messages to see if they are applicable. The most significant bit of a system message is always set to 1, just as for a channel message. The next three bits carry the system message code (111). No channel numbers are used with system messages, so the four least significant bits are free to carry the code that indicates the message type. This enables up to sixteen different system messages to be accommodated, but not all of the available codes are actually used at present. Figure 3 shows the way in which a system message header byte is coded.

System messages are divided into two broad categories, which are the system common and the system real-time messages. We will consider the system real-time messages first. These are mainly concerned with synchronising one MIDI unit to another, or possibly several slave units to the MIDI controller. This is only necessary when one or more of the slave units is being controlled by a built-in sequencer of some kind, rather than being controlled directly by the master unit. Many users prefer to have everything under the direct control of the master unit and this is certainly the method of control that I favour. However, if you wish to have (say) a MIDI drum machine under the control of its integral sequencer, it must be synchronised to the master sequencer. Table 4 lists the system real-time messages, and gives their four bit binary codes.

| Binary | Message Type |
|--------|----------------|
| 1000 | Clock Signal |
| 1001 | Undefined |
| 1010 | Start |
| 1011 | Continue |
| 1100 | Stop |
| 1101 | Undefined |
| 1110 | Active Sensing |
| 1111 | System Reset |

A conventional clock signal for a drum machine is a regular series of pulses and starting and stopping the clock signal starts

and stops the drum machine's internal sequencer. The MIDI clock message is roughly analogous to a conventional clock signal, but it is acceptable for it to be sent continuously. The slave sequencer is stopped and started using the stop and start messages. There is also a continue message, which has a different effect to the start type. A start message always starts the sequence 'from the top', even if the sequence had previously been stopped half way through. The continue message starts the sequence from its current position, wherever that may happen to be. The MIDI clock messages are sent at a rate of 24 per quarter note, incidentally. These are all simple single byte messages.

The song position pointer message is a system common type, but it is used in conjunction with the system real-time messages. It is used to move to any desired point in a sequence. The header byte is followed by two data bytes which together provide a 14 bit value. Figure 4 shows the way in which this system operates and it is basically the same method that is used to provide 14 bit values in pitch wheel change messages. This gives a range of 0 to 16383, but the smallest change this message can produce is six MIDI clocks (1/16th note). Using single clock resolution would presumably prevent the system from handling reasonably long sequences. To start a sequence at a given point, first a song position pointer message is used to take the sequencer to the appropriate point in the sequence, then a continue message is used to start the sequencer.

Common Sense

The active sensing is not used a great deal in practical MIDI systems and it is not implemented on many items of equipment. It has been included in the specifications of some recent instruments however, so it might be used to a greater extent in the future. It is just a simple fail-safe system that silences all the sound generator circuits of the slave instruments if a fault should occur, such as a broken cable. This avoids having slave instruments stuck with notes playing.

With active sensing in use, there must be a gap of no more than 300 milliseconds between one MIDI message and the next. If receiving units do not receive a MIDI message of some kind for more than 300 milliseconds, they shut down their sound generators. If a gap of more than this is about to occur, the controller will send an active sensing message to prevent the slave units from shutting down. The default is for no active sensing messages to be sent and this facility only comes into operation once the MIDI controller has sent the first active sensing message.

The system reset message does precisely what you would expect. It simply resets the slave instruments back to their default switch-on settings and switches off all notes. Many modern instruments have built-in memory circuits which remember the current settings when the instrument is switched off. When switched on again, instruments of this type simply carry on where they left off, rather than starting from a set of default settings. The correct effect for the system reset message with such instruments is debatable, but in most cases this message is simply not implemented. It is probably not a good idea to experiment with this message!

Next month
we will consider system exclusive messages
and General MIDI.

Computer controlled stepper motors

This month Jim Spence concludes his project using the ETI Forth Experimenter's Computer to accurately control the rotation of a stepper motor, with a look at the necessary software.

Last month we looked at how stepper motors work and the circuitry that would allow one to be accurately controlled by a computer. We could use any computer with a suitable parallel input/output card for testing this board, so long as all the inputs are TTL compatible. For the purposes of this article I will, however, assume you have built the Forth Experimenter's Computer published in the April 1994 and subsequent issues of ETI. This versatile single board computer with its in-built FORTH language is an ideal platform to test out such devices.

Table 1
WAVE STEP SEQUENCE

| STEP | A | B | C | D |
|------|---|---|---|---|
| 1 | 1 | 0 | 0 | 0 |
| 2 | 0 | 1 | 0 | 0 |
| 3 | 0 | 0 | 1 | 0 |
| 4 | 0 | 0 | 0 | 1 |

Table 2
TWO PHASE SEQUENCE

| STEP | A | B | C | D |
|------|---|---|---|---|
| 1 | 1 | 0 | 0 | 1 |
| 2 | 1 | 1 | 0 | 0 |
| 3 | 0 | 1 | 1 | 0 |
| 4 | 0 | 0 | 1 | 1 |

Table 3
Half STEP SEQUENCE

| STEP | A | B | C | D |
|------|---|---|---|---|
| 1 | 1 | 0 | 0 | 0 |
| 2 | 1 | 1 | 0 | 0 |
| 3 | 0 | 1 | 0 | 0 |
| 4 | 0 | 1 | 1 | 1 |
| 5 | 0 | 0 | 1 | 0 |
| 6 | 0 | 0 | 1 | 1 |
| 7 | 0 | 0 | 0 | 1 |
| 8 | 1 | 0 | 0 | 1 |

Step Sequence

The basic ways to step a stepper motor are 1 PHASE ON or WAVE, (Table 1) and TWO-PHASE-ON (Table 2), both of which are full step sequences. These can be combined to produce a half step sequence (Table 3).

The above sequences hold true for either unipolar or bipolar. The only difference is the way you connect the wires to IC1. Connect unipolar type motors, those with 5 and 6 wires, as described in the previous text, in the sequence the motor steps round. In the case of the bipolar motor, connect them as shown in Fig chop1. If it doesn't work then simply reverse one of the windings.

There are advantages and disadvantages of both single and half step methods of stepping (see table 4).

3 to the port. There is just one complication and that is the variable LAST-STEP. This simply remembers the last step so that each time the word is called it can advance to the next step. You should now be in a position to try out this word and see if it actually does advance the motor one half step at a time.

WARNING - watch the current, single stepping consumes far more current than fast stepping.

When you get bored watching the motor single step in response to you typing (stepf) you will notice that the motor is continuously drawing current. Also, if you try to turn it by hand, it appears to be 'locked'. To switch off the motor type 0 outs. You may like to define a word to do this, for example Box 3.

Box 4 shows how to implement

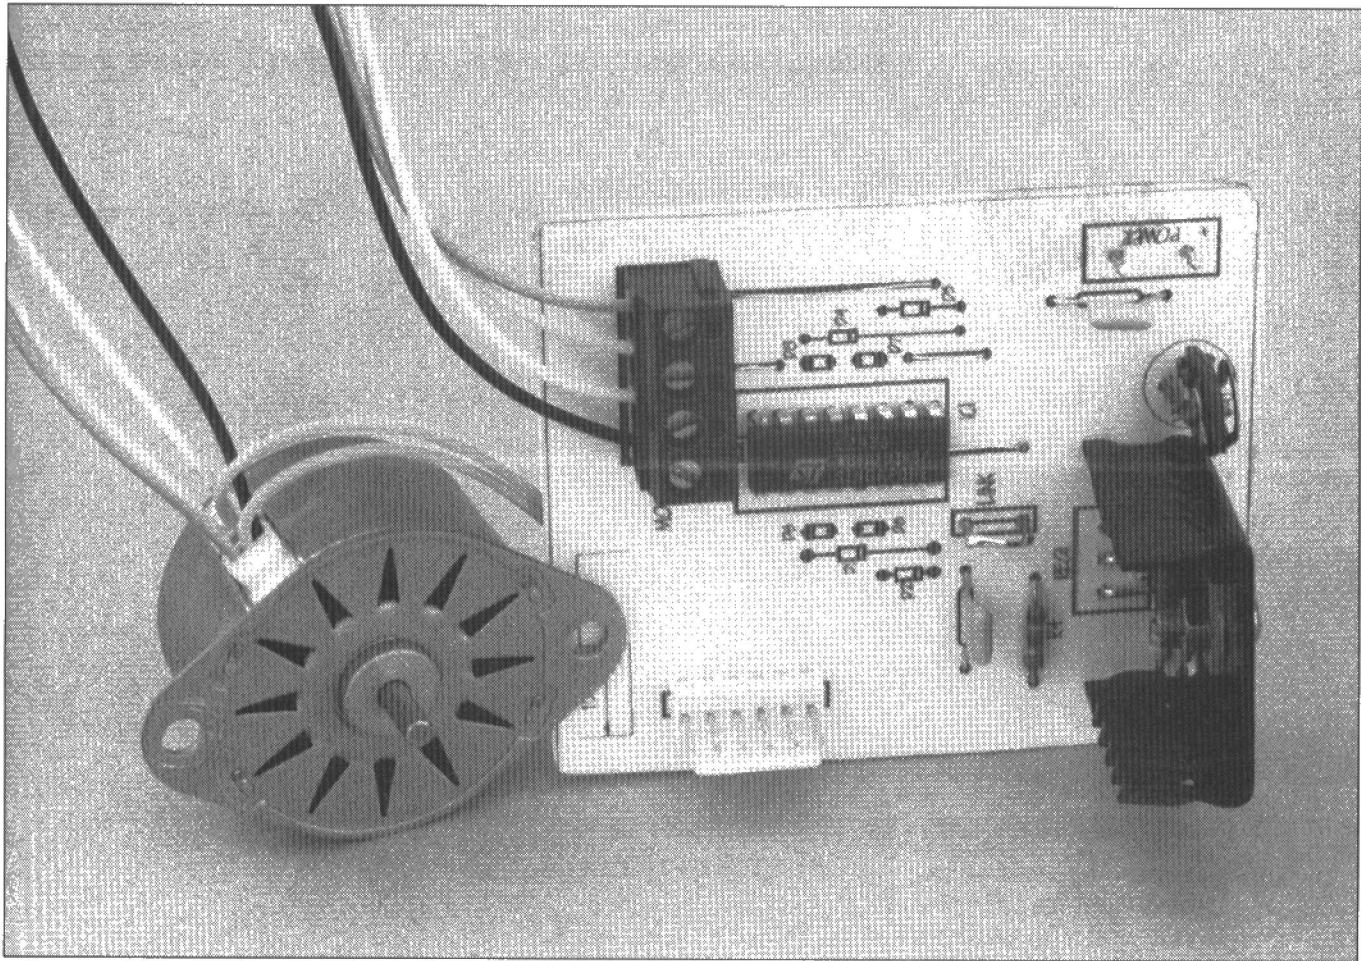
| Advantages | Single step | Disadvantages |
|---|-------------|--|
| • Equal torque for each step. | | • Resonance more likely. • Fewer, therefore larger steps per revolution. |
| Advantages | Half step | Disadvantages |
| • In the half step position motor only has half torque. • More steps per revolution. | | • Resonance much less likely. • Higher clock frequency required for same speed. |

The following software will be based on half step mode. This will be easy to change to full step if required.

First define a word for the output port we will be using (Box 1) Assuming that D0-D3 on the board is connected to D0-D3 on the computer, we will define a word called (stepf) 'step forward' which, each time it is called, will move the motor one half step forward (Box 2).

As you can see from this, it is simply outputting the steps as required in Table

continuous stepping and also has a little error checking. This will probably not work! Why? The simple reason is that it will try to step the motor too fast, unless by chance you are using a good motor. The stepping rate without a delay is about 1200 steps per second. Before we deal with this let's just go through the error checking. This is a word which is likely to be used interactively. It expects a value on the stack in the form of 20 step to step the motor 20 positions.



Now it's more than likely that from time to time you will forget to put something on the stack before typing step. The word depth leaves the number of items on the stack. The word 0= will leave a true value if the value on the stack is 0. So, if you type step without first putting anything onto the stack, depth will return 0 and 0= will return true. The 'if' will then activate and abort the program after displaying the error message.

Now to deal with the program working too fast. We obviously need a delay. This could be put in the do loop just after (or before) (stepf). As this delay is a fundamental requirement for the motor to work, it is probably best put in the lowest level word of all. This would be in the word outs (box 5).

We have also defined a new variable so that this can be adjusted to gauge the effects of the pull-in and pull-out. Don't forget that you must recompile all the words after outs so that the newly recompiled words will use the new version of outs. In other words, load them in again.

By now you should have enough words to make the motor whiz round to your heart's content. If you have a 48 step motor, half step mode would be 96

steps per one complete revolution. If you wanted it to turn 15 times you wouldn't need to get your calculator out, simply type 96 15 * step.

To backwards step, you simply reverse the sequence of steps, in our case from 1, 3, 6, 2, 4, etc. to 9, 8, c, 4, 6, etc. You can write another word called (stepb) to step backwards.

I mentioned earlier about ramping, increasing the speed of a motor after it has started. One solution is in boxes 6 and 7. Box 6 defines the essential variable required to store the ramp values. The idea is to step the motor 7 times at the slow rate, then step it another 7 times at an increased rate and then do the rest of the steps at the final rate. In practice, this works quite well. Of course 7 is a completely arbitrary number of steps, it's up to you to experiment with this in your application.

Improvements

There are lots of improvements which could be made. How about a variable which controls the direction of the step and perhaps another which controls full step or half step? If the motor is connected to something with a large mass, then consideration should be given to stepping the motor back just

before stopping. This would have a braking effect. Another idea which springs to mind is, instead of having the step sequence controlled by a 'case' statement, what about loading a byte with some value and rotating this to obtain the correct sequence?

Further Thoughts

As it stands, the circuit needs 4 wires to control the motor. It is tempting to consider reducing the number of control lines required. Perhaps this could be done by using the specialist circuitry as shown in Figure chop1. It would still need 4 lines for full control - 1) cw/ccw, 2) clock, 3) half/full and 4) enable. Depending on the application, you may get away with just direction and clock but without an enable this would leave the motor consuming power at stand still.

Advances are being made all the time in the design and manufacture of stepper motors. This article has concentrated on the types most likely to be found by the average amateur.

Acknowledgements

I would like to thank John Wood for his technical help on the finer points of driving stepper motors.

PARTS LIST

C1 L293 or L293D
REG1 L200CV
R1,R2,R3 1R 0.25W
R4 820R 0.25W
VR1 10K lin pre-set
C1,C2 0.1µF
D1-D8 BYW98-100 *
4 x PCB pins
1 x 6 way PCB connector
2 x PCB 2 way terminals
(Maplin JY92A)
* Diodes must have a Tr of
less than 200ns, not needed if
type L293 is used.

Basic out put word

```
: outs
\ Stepper output port
0 p!
;
```

Step forward Word

```
hex
: (stepf)      ( -- )
\ Step forward
last-step @
\ do next forward step
case
  1 of 3 dup outs
endof
  3 of 2 dup outs
endof
  2 of 6 dup outs
endof
  6 of 4 dup outs
endof
  4 of c dup outs
endof
  c of 8 dup outs
endof
  8 of 9 dup outs
endof
  9 of 1 dup outs
endof
  ." Error in
variable last-step "
  endcase
last-step !
;
```

Word used to switch off motor

```
: n 0 outs ;
```

Step with error checking

```
: step      ( x -- )
depth 0=      \
check something on
stack
if
  cr ." Number of
steps not specified"
  abort
else
  0 do      \
begin loop for number
of steps
  (stepf)      \
do one step
  loop      \
continue until all done
endif
  0 outs      \
switch motor off
  ;
```

Outs with delay

```
decimal
variable del-value
300 del-value !
: delay
  del-value @ 0 do loop
  ;
: outs
\ Stepper output port
0 p!
delay
;
```

Defining ramp variables

```
decimal
work in decimal
variable ramp1      \
define 3 ramp variable
variable ramp2      \
for 3 ramp stages
variable ramp3
300 ramp1 !      \
initialise variables
100 ramp2 !
50 ramp3 !
```

One method of ramping

```
decimal
: rstep      ( x -- )
)  \ x is the number
of steps
ramp1 @ del-value !
adjust delay to first
ramp value
dup
number of steps
0
do
step
step motor 1
1-
reduce number of steps
i 7 =
compare index with 7
?leave
leave if index = 7
loop
else continue to loop
ramp2 @ del-value !
adjust delay to first
ramp value
dup
number of steps
0
do
step
step motor 1
1-
reduce number of steps
i 7 =
compare index with 7
?leave
leave if index = 7
loop
else continue to loop
ramp3 @ del-value !
adjust delay to first
ramp value
0
?do
only do loop if any
steps left
step
step motor 1
loop
else continue to loop
last-step @ outs
delay  \ stop motor
0 outs
turn off motor
;
```

Practically Speaking

by Terry Balbirnie

Mains Free Zone

Over the past two months, we have looked at the setting-up of a workshop inside the house, garage or garden shed. It may be necessary to avoid a mains supply either for reasons of safety (perhaps because the location is damp) or because it is too expensive to install on account of the distances involved.

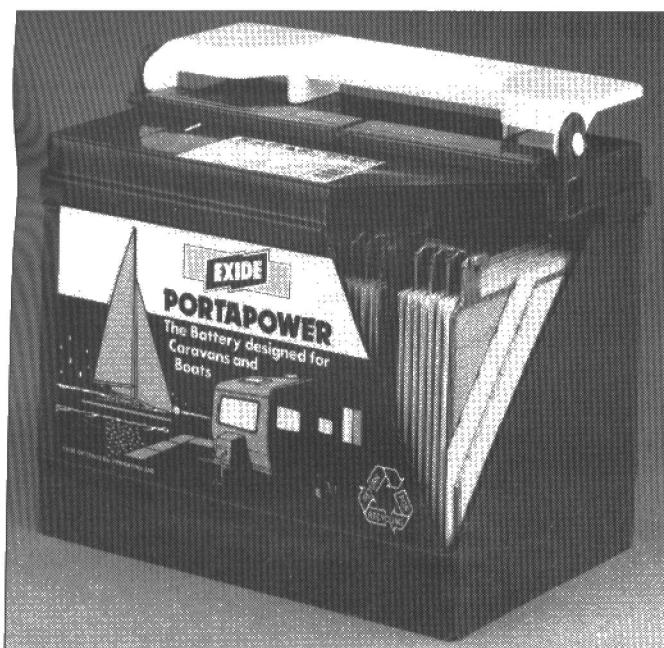
Do not despair! It is perfectly possible to operate a mains-free workshop, but an alternative power supply is needed.

Battery Power

A 12V car-type battery can be used as a supply for soldering, lighting and testing experimental circuits. However, where the battery is regularly run down to a low state of charge (so-called cyclic use) you will need a 'leisure' battery. These are designed chiefly for caravan and boat users and will withstand the type of cyclic use which a standard car battery will not. The illustration shows one from the Exide Portapower range. These are available in three capacities and are fitted with a carrying handle that makes them easy to move around.

A leisure battery will provide years of service, so long as it is not abused by overcharging, short-circuiting or being left in a discharged state. If any lead-acid battery is left discharged it will soon be ruined. The battery should be charged promptly when there is any sign of the output voltage falling to less than about 11V on load (e.g. by lights dimming slightly). A narrow-scale voltmeter covering the range 10 to 15V will be described as a constructional project in a later issue and this will be found useful for battery checking. It is a good rule to maintain the battery in as high a stage of charge as possible, as even slight abuse will lessen the capacity of the battery and it may never recover.

Wherever the battery is used, standing idle or on charge, it is essential to keep it well away from sparks, naked flames,



etc., and always site it in a well-ventilated area. There may be some hydrogen gas around and lighting this will cause an explosion. Remember also that a battery will not hold its charge indefinitely and during periods of non-use, it should be charged every three months. At the same time, the electrolyte level should be checked according to the manufacturer's instructions. This will also be a good time to lightly smear the terminals with Vaseline to prevent corrosion. The battery is unlikely to require topping up unless it has been carelessly charged or is nearing the end of its useful life. If the level is low, top up with distilled or de-ionised water - never tap water.

I'm On Charge

Leisure batteries must be charged with respect. To avoid any chance of sparking and possible ignition of hydrogen gas, the battery should be connected before the charger is switched on. Similarly, it should be switched off before disconnecting the battery. If a car-type battery charger is used, its terminal voltage should first be measured. This should not exceed 14.4V. Also, the battery should not be charged for more than 12 hours continuously. Note that the maximum charge rate must never be exceeded. A constant-voltage charger with current limiting will be described as a constructional article in a future issue.

The Labcraft TP2 MkVI (available from caravan dealers) is a complete battery management system. It has a compartment into which the battery is placed. It may be fitted with one of two chargers - 4A or a 10A (maximum) switched mode type. Both are fully regulated so that the battery cannot be overcharged. There is a fused output socket and a voltmeter which acts as a state-of-charge indicator. The unit has a carrying strap so that it may be conveniently carried back to the house for charging. The chargers are also available separately for free-standing use.

My Resistance Is Low

When using any lead-acid battery as a power supply, remember that it has an exceptionally low internal resistance. This means that if it is short-circuited by applying, say, a piece of thick wire to its terminals, an enormous current will flow - perhaps several hundred amps! This will generate a lot of heat, probably melt the wire, possibly burn the skin and pose a fire risk. It will also damage the battery itself.

For these reasons, it is essential to provide a fuse in the positive feed wire close to the battery. A 5A or 10A fuse in an enclosed car-radio type fuse holder will be adequate for workshop purposes. This should be used with proper battery connectors. Do not use crocodile clips on the terminals. With a unit such as the TP2, of course, there are already connectors, fuse, etc., fitted. From here, the appliances may be plugged into a distribution board made by mounting 2-pin chassis sockets in a plastic box. The American flat-pin type would be suitable.

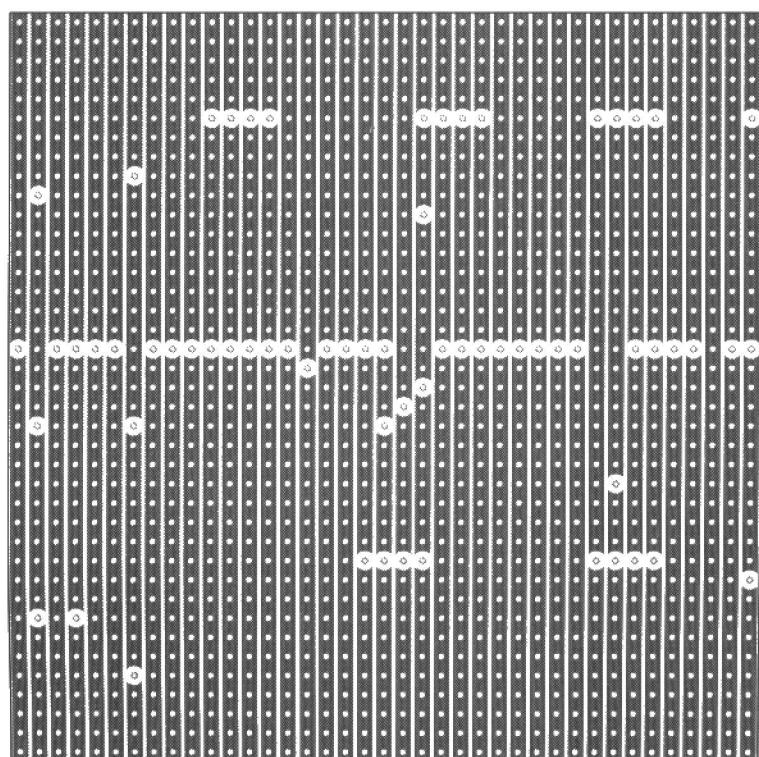
Next month

we shall continue by looking at the choice of battery capacity to suit individual needs and look at the type of equipment available for use in a mains-free workshop.

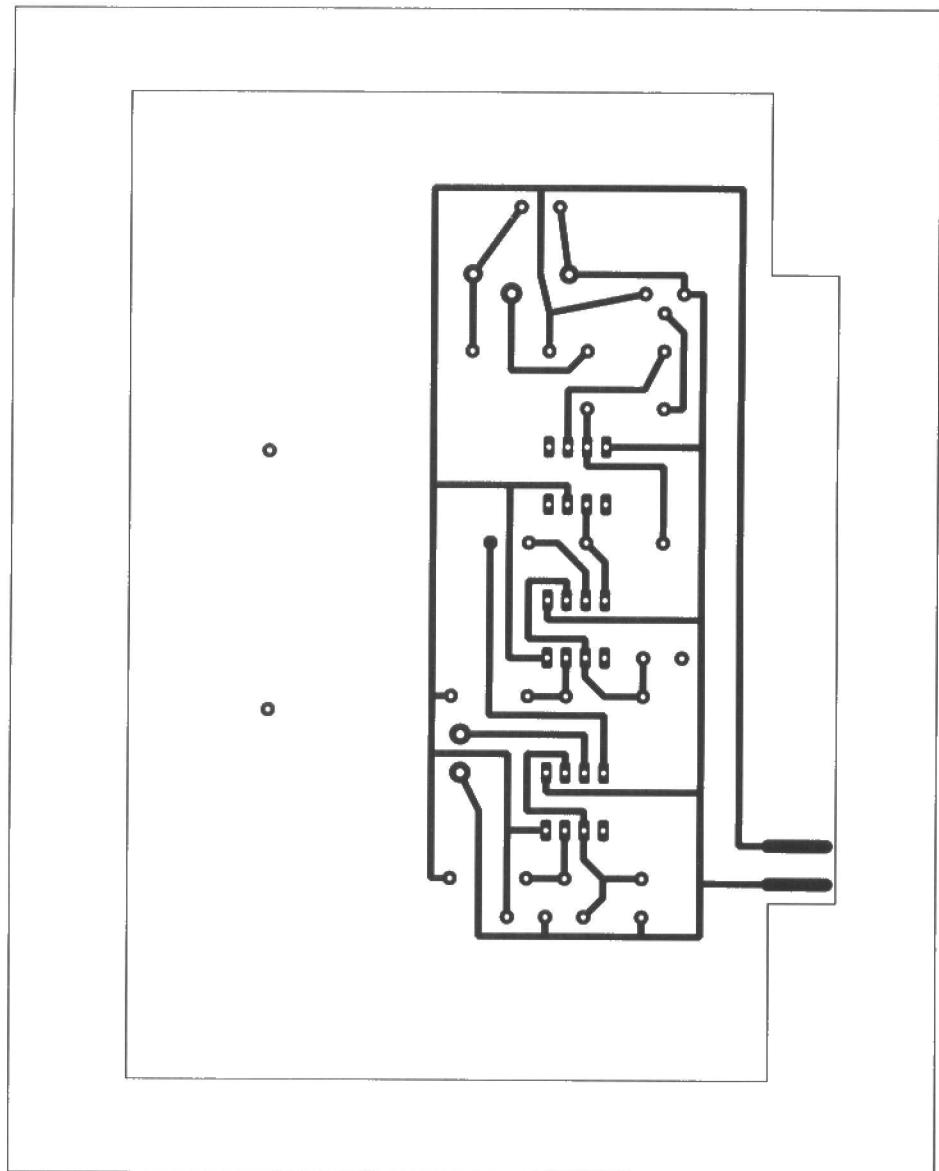


ELECTRONICS
TODAY INTERNATIONAL

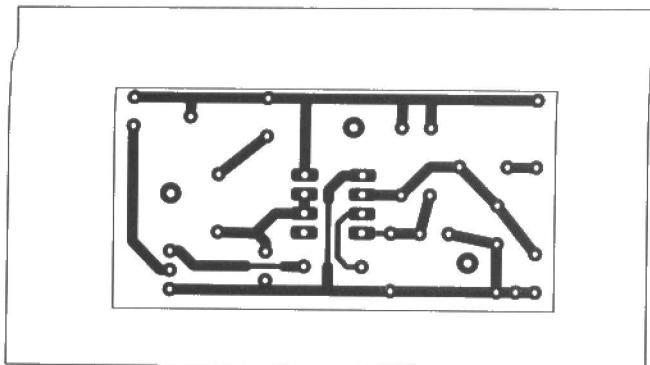
Foils for this issue



A Touch of Magnetism



PC Overheating Alarm



Tow Bar Alarm

Open Forum

It is now twelve months since I took over as editor of ETI, with the clear brief to make it the most popular and widely read practical electronics magazine in the country. No easy task, but one that, with a lot of fumbling and experimentation, is on the verge of being achieved.

If you are a regular reader, you will have noticed changes in the magazine. There are more pages and they are printed on better paper, so hopefully you will find the diagrams and photographs clearer and the text easier to read. These changes have, we think, made the magazine look better and more pleasant to read.

The editorial has also changed slightly. Hopefully, ETI is still providing you with the type of projects which you have traditionally come to expect, but I am also attempting to broaden the appeal of the magazine so that it will be read by people who are not necessarily committed electronic project builders and also by those who find the traditional projects technologically conservative.

To achieve these aims, I have added electronically oriented feature articles which should appeal to anyone with a general interest in, and understanding of, technology. I am also trying to make the magazine of interest to the enormous number of PC users who have an interest in computer technology, hence the PC Clinic series and projects to make useful PC add-ons.

For the technologically adventurous, there are the high quality projects, such as the world's first Transputer project in an electronics magazine, or the FORTH experimenter's computer.

These projects put the builder at the forefront of technology and will be of equal interest to commercial and academic readers as they are to our traditional hobbyist readers.

The next stage in the editorial development of ETI is to encourage some positive feedback from you, the reader. One of the major problems with any magazine is that one only receives the criticism for mistakes that are made, never any comment about things that are going right. In other words, I need to know what you think of each article and whether there is anything you would especially like to

see in the magazine over the coming months.

We have already started to gather this type of information. The observant amongst you will have noticed in previous issues a little piece tucked away on the letters page entitled Feedback. The initial trials have been very successful and I would now like to take this opportunity to encourage more readers to respond to what will be a regular monthly questionnaire.

It will be extremely helpful to me if, just once in a while, readers will take a few minutes and write down on the back of a postcard the ratings which they would award to each article in the current issue. Ratings should vary between 1 and 10, with 1 being poor and 10 being brilliant, and if there is something that you would like to see covered in ETI, why not make a brief note of it after the ratings list.

Only by knowing what you, the readers, want can I endeavour to make ETI better. I don't want ETI to be simply the best electronics magazine in Britain but the best in the world - together we can do it!

Thanking you for your support.

Nick Hampshire, Editor

Feedback September 1994.

- A - The Electronic Nose
- B - PC Overheating Alarm
- C - PC Clinic
- D - Computer Controlled Stepper Motors
- E - An Introduction to MIDI
- F - Constant Current Transistor Tester
- G - Tow Bar Alarm
- H - Fluxgate Magnetometer
- I - ETI Transputer Board
- J - Practically Speaking

Just write the article letter followed by your score for that article and send it clearly marked to Feedback Box September 94, ETI, Argus House, Boundary Way, Hemel Hempstead, Herts. HP2 7ST.

To add an extra incentive, all replies received before August 30th 1994 will go into a draw and the winner will receive a couple of handy cases for ETI projects, so don't forget to put your name and address on the card.

Next Month...

Next month we will start the construction of ETI's world beating single board Transputer project, from designers Mark Robinson and Andy Papageorgiou. We will also be starting the construction of Stephen Smith's Power On Self Test card, that will enable any PC owner to accurately locate faults in the hardware. For electronic music enthusiasts there is a Hall Effect Swell Pedal from Robert Penfold, as well as the next part of his MIDI tutorial.

For fans the FORTH Experimenter's Computer, there is a project to use this board to decode the National Physical Laboratory Rugby 60kHz time signal. There is a handy 12V lead acid battery charger from Terry Balbirnie and we will also be concluding Keith Garwell's magnetometer project.

The main feature will be a look at the global optical 'Information Super Highway', that is already destined to change the way we work, the way we shop and the kind of entertainment that comes into our homes.

Don't miss October's ETI. Why not place an order with your newsagent now?



EDITORIAL

Editor Nick Hampshire

Sub Editor Jim Bluck

CREATIVE

Designer Nadia Ahdout

Technical Illustration John Puczynski

Photography Manny Cefal

ADVERTISEMENT SALES

Display Sales

Hellen Taylor

Advertisement Copy Control

Marie Quilter

Classified Sales

James Gale

MANAGEMENT

Managing Director

Terry Pattisson

Circulation & Promotions Manager

Debra Stupple

Production Manager

Tony Dowdeswell

Group Editor

Stuart Cooke

Group Advertisement Manager

Claire Jenkinson



ISSN
0142-7229

ETI is normally published on the first Friday in the month preceding the cover date. The contents of this publication including all articles, plans, drawings and programs and all copyright and all other intellectual property rights therein belong to Argus Specialist Publications. All rights conferred by the Law of Copyright and other intellectual property rights and by virtue of international copyright conventions are specifically reserved to Argus Specialist Publications and reproduction requires the prior written consent of the company c1990 Argus Specialist Publications. All reasonable care is taken in the preparation of the magazine contents, but the publishers cannot be held legally responsible for errors. Where mistakes do occur, a correction will normally be published as soon as possible afterwards. All prices and data contained in advertisements are accepted by us in good faith as correct at the time of going to press. Neither the advertisers nor the publishers can be held responsible, however, for any variations affecting price or availability which may occur after the publication has closed for press.

Subscription rates-UK £23.40 Europe £30.30 Sterling Overseas £31.50 US Dollar Overseas \$52.00

Published by Argus Specialist Publications, Argus House, Boundary Way, Hemel Hempstead HP2 7ST. Telephone (0442) 66551. UK newstrade distribution by SM Distribution Ltd, 6 Leigham Court Road, London SW16 2PG. Telephone 081-867 8111. Overseas and non-newstrade sales by Magazine Sales Department, Argus House, Boundary Way, Hemel Hempstead, HP2 7ST. Telephone (0442) 66551. Subscriptions by Argus Subscription Services, ETI, Queensway House, 2 Queensway, Redhill, Surrey RH1 1QS. Telephone (0737) 788611. US subscriptions by Wise Owl Worldwide Publications, 4314 West 238th Street, Torrance, CA90505 USA. For Visa/Mastercard orders in USA - Telephone (310) 375 6258 Fax (310) 375 0548. Pacific Time: 9am-9pm Weekdays. 10am-6pm Weekends. Typesetting and origination by Ebony, Linkerd, Cornwall. Printed by Wiltshire Ltd. Bristol.



**Argus House, Boundary Way,
Hemel Hempstead HP2 7ST
Telephone (0442) 66551
Fax (0442) 66998**

NEW

FULL COLOUR GUIDE TO ELECTRONIC PRODUCTS

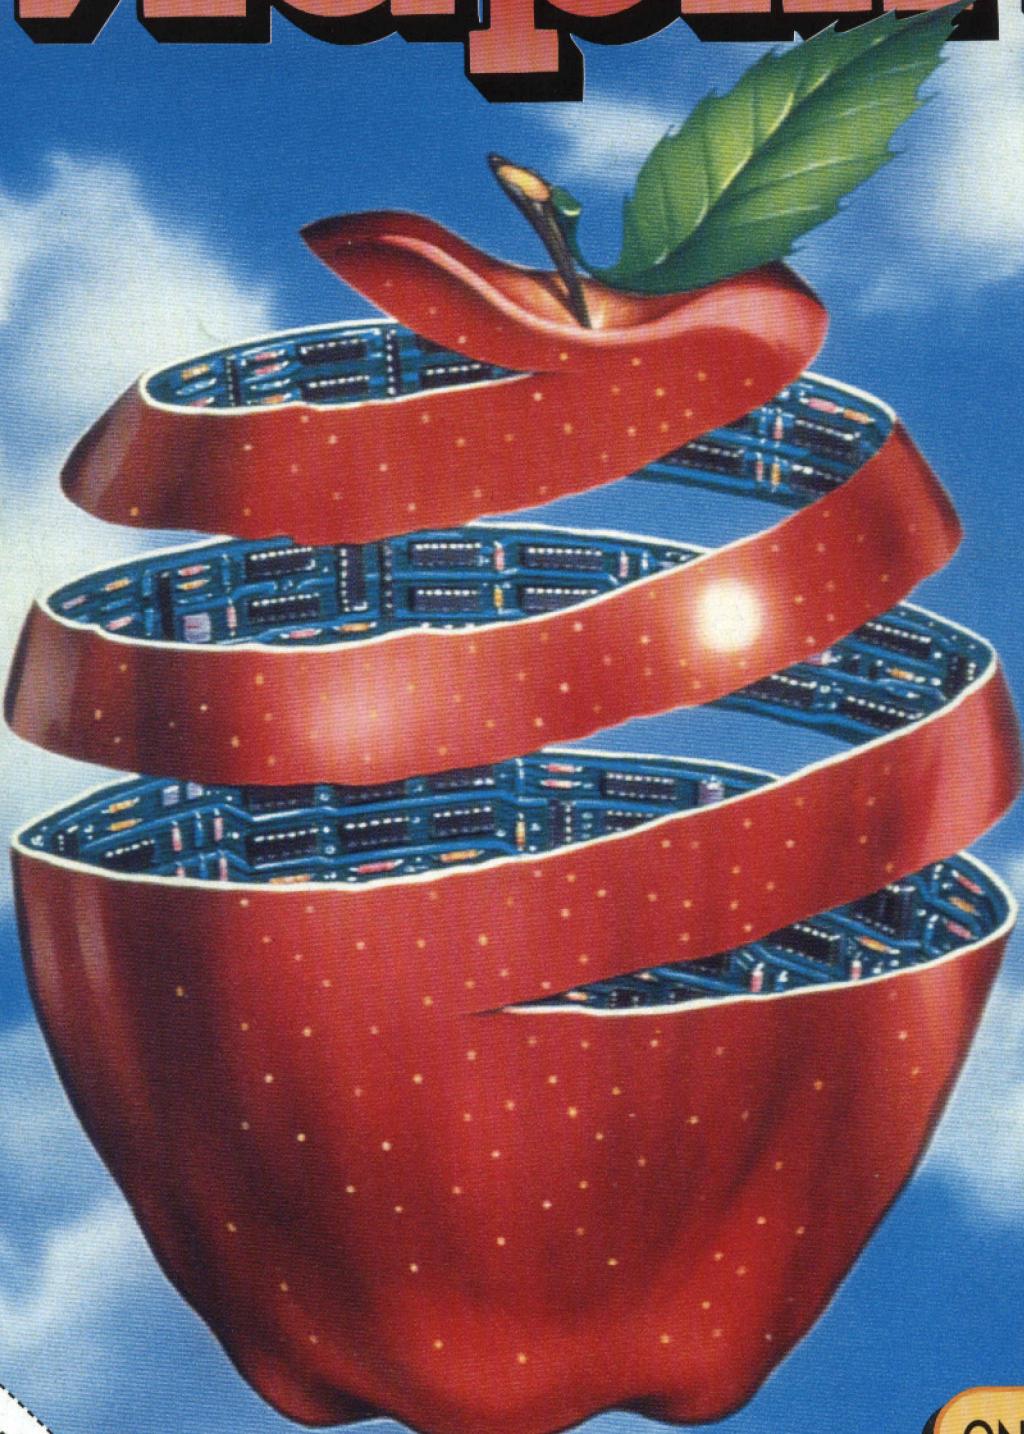
Sept 1994–Aug 1995

Maplin



BS 5750
Part 2 1987

Level B:
Quality Assurance
RS12750



Order your copy of the New MAPLIN Catalogue on sale 2nd September!
Pick up a copy from any branch of WHSMITH, branches of John Menzies in Scotland ONLY, Eason & Son in N. Ireland ONLY, and Maplin stores nationwide for just £3.45 or post this coupon now to receive your copy for just £3.95 inc. p&p. If you live outside the U.K. send £7.30 or 18 IRCs for Airmail in Europe (including Republic of Ireland); £6.50 or 16 IRCs for surface mail outside Europe; or £12.30 or 30 IRCs for Airmail outside Europe. I enclose £3.95/£7.30/£6.50/£12.30 (delete as applicable).

Name
Address
Post Code
Send to Maplin Electronics,
P.O. Box 3, Royleigh,
Essex SS6 8LR
EN9 9S

Over 800 colour packed pages
with hundreds of Brand New
Products at Super Low Prices.



Available from all branches of WHSMITH, John Menzies in Scotland ONLY, Eason & Son in N. Ireland ONLY, and Maplin stores nationwide. The Maplin Electronics 1995 Catalogue – OUT OF THIS WORLD!

ON SALE FROM
2ND SEPTEMBER
ONLY
£3.45